Session 3 Overview: *Ultra-High-Speed Wireline Transceivers and Energy-Efficient Links*

WIRELINE SUBCOMMITTEE

Session Chair: *Hyeon-Min Bae, KAIST, Daejeon, Korea*

Session Co-Chair: *Ajith Amerasekera, Texas Instruments, Dallas, TX*

Smart phones and their social apps are key drivers for the growth of the internet and, more generally, of big data infrastructure. The ever-increasing demands placed on this infrastructure, in turn, spurs the insatiable need for data communication bandwidth between chips. In this context, wireline transceivers that push the limits imposed by process technology – in terms of data rate, energy efficiency, and area – are extremely critical. This session introduces 2 ADC-based receivers operating at or beyond 25Gb/s for multi-standard support. It continues with a presentation of a 25Gb/s equalizer for cable-dominated channels with up to 50dB of half-baud loss. Two papers addressing different approaches for high-data-rate, short-reach communication are presented next, including an NRZ/PAM-4 Ethernet ADC-based transceiver and a 56Gb/s per lane NRZ transceiver. Finally, the session concludes with two papers describing critical components for I/O solutions at per-lane rates at or beyond 45Gb/s, including a power-efficient 45Gb/s PAM-4 transmitter and a 64Gb/s NRZ transmitter for CDE-56-VSR/MR applications, the latter implemented in 16nm finFET.

1:30 PM

3.1 A 25Gb/s ADC-Based Serial Line Receiver in 32nm CMOS SOI

S. Rylov, IBM T. J. Watson Reseach Center, Yorktown Heights, NY

In Paper 3.1, IBM presents a 25Gb/s ADC-based serial line receiver in 32nm CMOS SOI. The receiver uses a ¼ rate 5b flash ADC with on-chip calibration followed by an interleaved 8-tap FFE and 8-tap DFE backend and digital CDR. The transceiver occupies 0.39mm^2 and compensates a reflective 40dB loss channel while consuming 453mW.

2:00 PM

3.2 A 320mW 32Gb/s 8b ADC-Based PAM-4 Analog Front-End with Programmable Gain Control and Analog Peaking in 28nm CMOS

D. Cui, Broadcom, Irvine, CA

In Paper 3.2, Broadcom presents a 32Gb/s 8b ADC-based PAM-4 receiver in 28nm CMOS. The ADC achieves an ENOB of 6.4 and 5.85 at DC and Nyquist, respectively. The entire receiver, occupying 0.89mm^2, features a continuous-time linear equalizer with a 7dB peaking gain, and compensates for channel loss of more than 50dB at 16GHz while consuming 320mW.

 978-1-4673-9467-3/16/$31.00

2:30 PM

3.3 A 25Gb/s Multistandard Serial Link Transceiver for 50dB-Loss Copper Cable in 28nm CMOS

T. Norimatsu, Hitachi, Tokyo, Japan

In Paper 3.3, Hitachi presents a 25Gb/s multi-standard serial link transceiver in 28nm CMOS. The receiver, using a 20dB CTLE and a 14-tap DFE, achieves 3mV input sensitivity by using sub-mV dynamic DC-offset cancellation and DFE tap-bias control schemes. The transceiver, occupying 5.5x4.6mm^2, achieves BER < 10^{-12} over a 51dB-loss 5m AWG channel at 25Gb/s and consumes 403mW from 0.9 and 1.5V power supplies.

3:15 PM

3.4 A 40/50/100Gb/s PAM-4 Ethernet Transceiver in 28nm CMOS

K. Gopalakrishnan, Inphi, Santa Clara, CA

In Paper 3.4, Inphi describes a 40/50/100Gb/s NRZ/PAM-4 Ethernet transceiver in 28nm CMOS. The receiver integrates two SAR ADCs with ENOB of 4.8, a DSP for calibration, and an FEC encoder. The transmitter includes a FEC decoder and supports both NRZ and PAM-4 signaling with a maximum swing of 1.4Vpp. The transceiver, which includes an internal supply voltage regulator, occupies 6.3x4.9mm^2 and achieves less than 200fs output jitter while consuming 2.4W from 0.9V and 1.2V power supplies.

3:45 PM

3.5 A 56Gb/s NRZ Electrical 247mW/lane Serial Link Transceiver in 28nm CMOS

T. Shibasaki, Fujitsu Laboratories, Kawasaki, Japan

In Paper 3.5, Fujitsu presents a 56Gb/s/lane electrical NRZ transceiver in 28nm CMOS for high-data-rate applications, using baud-rate sampling to minimize the power consumption. The design features a 2-tap feed-forward equalizer and a continuous-time linear equalizer, followed by a 1-tap speculative decision-feedback equalizer. The receiver compensates for 18.4dB half-baud loss. The transceiver occupies 1.4mm^2 per two lanes and consumes 247mW/lane from a 0.96V supply.

4:15 PM

3.6 A 45Gb/s PAM-4 Transmitter Delivering 1.3V_{ppd} Output Swing with 1V Supply in 28nm CMOS FDSOI

M. Bassi, University of Pavia, Pavia, Italy

In Paper 3.6, researchers from the University of Pavia present a 45Gb/s PAM-4 transmitter able to deliver 1.3Vppd output swing; the design is implemented in 28nm fully-depleted SOI CMOS and is intended to address next generation CEI-56G standards for high speed links. The design occupies an area of 0.28mm^2 and consumes 120mA from a 1V supply.

4:45 PM

3.7 A 40-to-64Gb/s NRZ Transmitter with Supply-Regulated Front-End in 16nm FinFET

Y. Frans, Xilinx, San Jose, CA

In Paper 3.7, Xilinx presents a 3-tap 64Gb/s NRZ transmitter using a quad-rate architecture for short reach electrical links implemented in a 16nm FinFET technology. The design, which is applicable to next generation standards such as CEI-56-VSR/MR, uses power efficient techniques that take into consideration the 16nm FinFET device properties. The transmitter achieves 800mVppd output swing with 150fs random jitter while consuming 225mW at 64Gb/s and occupying an area of 0.32mm^2.

3.1 A 25Gb/s ADC-Based Serial Line Receiver in 32nm CMOS SOI

Sergey Rylov[1], Troy Beukema[1], Zeynep Toprak-Deniz[1], Thomas Toifl[2], Yong Liu[3], Ankur Agrawal[1], Peter Buchmann[2], Alexander Rylyakov[4], Michael Beakes[1], Benjamin Parker[1], Mounir Meghelli[1]

[1]IBM T. J. Watson Reseach Center, Yorktown Heights, NY,
[2]IBM Zurich Research Laboratory, Rüschlikon, Switzerland,
[3]now with Shanghai Jiao Tong University, Shanghai, China,
[4]now with Coriant Advanced Technology Group, New York, NY

As CMOS devices continue to scale down in voltage and area, digital-based high-speed serial I/Os [1] become increasingly competitive with analog-based designs [2,3]. In addition to offering the PVT-independent performance of digital functions and superior power and area scaling to future technology nodes, digital-based I/Os can support advanced line modulation techniques that will become necessary as long-reach electrical channel data rates scale to 56Gb/s and beyond. The key enablers of a digital receiver are power and area efficient analog to digital conversion (ADC) and digital channel equalization. This paper describes the design of a 25Gb/s 2-level digital serial line receiver including a ¼-rate 5b flash ADC, an 8-tap feed-forward equalizer (FFE), an 8-tap decision-feedback equalizer (DFE), and a baud-rate clock and data recovery circuit (CDR). The receiver features a flash ADC, which employs a new power and area efficient slicer design capable of achieving high-precision (~1mV) threshold accuracy with an associated on-chip calibration system. The 32nm SOI CMOS receiver achieves error-free operation with margin on a reflective transmission-line channel with 40dB half-baud loss.

The receiver block diagram is shown in Fig. 3.1.1. The input uses AC-coupling (ACC) to a variable-gain amplifier (VGA) with a feed-forward DC restore network [2]. The VGA drives a two-stage peaking amplifier (CTE) boosting high frequency gain by ~12dB at 12.5GHz [3]. The CTE drives a 25GS/s ¼-rate 5b ADC organized as two half-ADCs, HADC-I and HADC-Q, driven by differential clocks CI (0°, 180°) and CQ (90°, 270°) derived from the phase rotator (PR) output with a divide-by-2 circuit. The aligned ADC output (40b at 1/8th-rate) is sent to a logic back-end containing a CDR block, a digital signal processing (DSP) block, and a state machine (Cal Engine) for offset calibration of the ADC comparators. The Cal Engine employs a 12b differential DAC (CALDAC) made of a pair of 2nd-order $\Delta\Sigma$ 1b DACs followed by 3rd-order passive RC filters.

Figure 3.1.2 shows a block diagram of one 5b HADC that uses a ping-pong flash architecture. Its two sets of 30 comparators form a stack of 15 2×2 data quads (DAT QUAD) sharing a single resistive ladder. Each comparator receives an input sample D from the ADC master-slave sampler and a reference voltage R from the ladder. The master sampler stage consists of an NFET source follower (NSF) driving two sampling capacitors via bootstrapped switches on complementary clock phases (A and B), while the slave stage consists of PFET source followers (PSF) driving bootstrapped switches embedded into individual comparators. The slave stage holds charge passively on comparator inputs connected by a floating strap S to minimize kickback-induced offsets.

Both NSF and PSF employ the same push-pull topology (shown in Fig. 3.1.2 for NSF) targeting slightly over unity gain to reduce the CTE output signal swing and improve its operating linearity. To operate the NSF in push-pull mode, the gate voltage of the tail device M2 in its primary branch (M1, M2) is inverted with respect to the gate of M1. This is primarily achieved via current folding of an extra NSF branch (M3, R2) via PFETs M0, M6. Each ADC stack includes an additional calibration quad (CAL QUAD). The CAL QUAD contains two shadow comparators used to substitute any of the 60 data comparators for offset calibration. The output of all data and shadow comparators is brought to the 1/8th-rate clock domain (C8) and passed to a logic block T2B for bubble correction, metastability hardening, thermometer-to-binary conversion, and shadow substitution. The Cal I/O logic block connects the Cal Engine with the comparator selected for calibration.

Figure 3.1.3 shows a block diagram of one time-interleaved comparator pair (A, B). Each comparator uses an integrating summer to subtract reference R from its data input M, which is derived from the input ports D, S and V via a 3:1 multiplexer serving also as a slave sampler switch. Since each summer in a pair integrates for only half of the time, their current tails are time-shared to save power. These currents are calibrated with a replica circuit as described in [3]. The multiplexer is controlled with a logic signal, CAL. In regular mode (CAL=0), the summer input M connects to the floating strap S for passive hold and connects for ½ a cycle to the sampled data port D; in calibration mode (CAL=1), it connects to the precise voltage reference V from the CALDAC. The integrating summer is loaded with a DCVS latch (slicer) in parallel with a 9b passive capacitive DAC (CDAC) that cancels comparator offset. The slicer output is applied to a demultiplexer cell that sends 1/8th-rate outputs Z0, Z1 to the logic block T2H when CAL=0 and a low-speed output ZCAL to the Cal Engine when CAL=1.

The high-resolution CDAC uses a switched-capacitor array as shown in Fig. 3.1.4. Its entire coarse 6b-section (except for its LSB) is made of identical single-finger PFETs serving as capacitors or as switches. The switches steer blocks of 0, 1, or 3 finger capacitors between P and N CDAC terminals. Note that an "on" switch counts as an extra capacitive finger bringing the block total to 1, 2, or 4 fingers. The coarse LSB is implemented with half-sized PFET fingers, whereas the 3b fine section uses those half-size PFET switches in series with smaller wire capacitors sized for a 0.5b coarse/fine range overlap. The measured CDAC transfer curve in Fig. 3.1.4 shows its sub-mV resolution over a range of +/-90mV (referred to the summer input).

The aligned 1/8th-rate ADC digital output is sent to two separate paths: (1) a primary DSP processing path to apply FFE8/DFE8 equalization and provide a sample "snapshot" function, and (2) a second low-latency FFE2/DFE1 path to implement a baud-rate CDR function. As shown in Fig. 3.1.5, the DSP consists of an FFE8 interleaved in blocks of 4 with a 4+4 tap DFE to provide a combined ISI coverage range of 2 pre-cursors and 12 post-cursors. The FFE design is a power and speed-optimized version of the distributed-arithmetic design [4] with added digital gain control (×1, ×2, ×4). The DFE, shown in detail in Fig. 3.1.5, features taps realized by look-ahead (taps 1 to 4) and direct feedback (taps 9 to 12), while carry rippling is avoided in the adder by feeding the outputs of a carry-save adder (CSA) stage directly into an 8b comparator [5].

The receiver functionality has been fully tested. The on-chip ADC threshold calibration system was found to produce higher-than-expected DNL (~3mV RMS or ~1/5 of a 15mV LSB vs. sub-mV CDAC resolution). Its likely cause is the CALDAC output impedance being not low enough to hold the slicer input port V stable in calibration mode. To take full advantage of the high-precision ADC threshold adjustment, an off-chip software-based ADC calibration was used to improve calibrated DNL to ~1mV RMS, allowing the ADC to achieve ~4.5b low-frequency ENOB and ~4b high frequency ENOB (Fig. 3.1.6). To test the complete receiver line equalization performance, a 25.6Gb/s PRBS23 data source was applied via a ~40dB-loss reflective transmission line channel. With equalization applied and CDR active, the receiver achieves an effective 18% UI (~8ps) horizontal eye margin at 1E-12 BER as shown in Fig. 3.1.6. A chip photo of the test site is shown in Fig. 3.1.7. The receiver operating at 25.6Gb/s dissipates 410mW from a 1V supply and 43mW from a nominal 0.7V supply (used in FFE only), for a total of 453mW.

Acknowledgement:
The authors thank M. Sanduleanu, JO Plouchart, J. Proesel, S. Dhawan, S. Kim, M. Kuppannan and A. Sahasrabuddhe for their contributions.

References:
[1] B. Zhang et al., "A 195mW / 55mW Dual-Path Receiver AFE for Multistandard 8.5-to-11.5 Gb/s Serial Links in 40nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 33-34, Feb. 2013.
[2] G. Gangasani et al., "A 32 Gb/s Backplane Transceiver With On-Chip AC-Coupling and Low Latency CDR in 32 nm SOI CMOS Technology", *IEEE J. Solid-State Circuits*, vol. 49, no. 11, pp. 2474-2489, Nov. 2014.
[3] J. Bulzacchelli et al., "A 28Gb/s 4-Tap FFE/15-Tap DFE Serial Link Transceiver in 32nm SOI CMOS Technology", *ISSCC Dig. Tech. Papers*, pp. 324-325, Feb. 2012.
[4] T. Toifl et al., "A 3.5 pJ/bit 8-tap-feed-forward 8-tap-decision feedback digital equalizer for 16Gb/s I/Os", *European Solid-State Circuits Conf.*, pp. 455-458, Sept. 2014.
[5] G. Linnenberg, "Two-sum comparator", US patent, 5,144,577, published Sept. 1992.

 978-1-4673-9467-3/16/$31.00

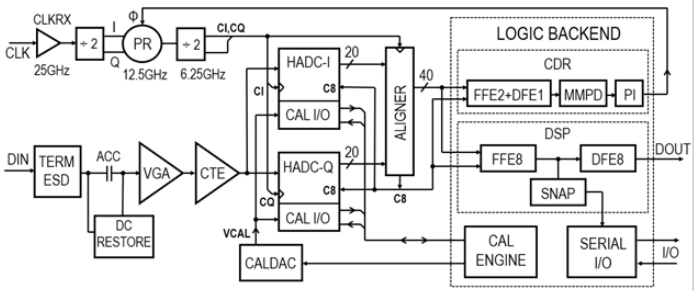

Figure 3.1.1: Block diagram of the receiver.

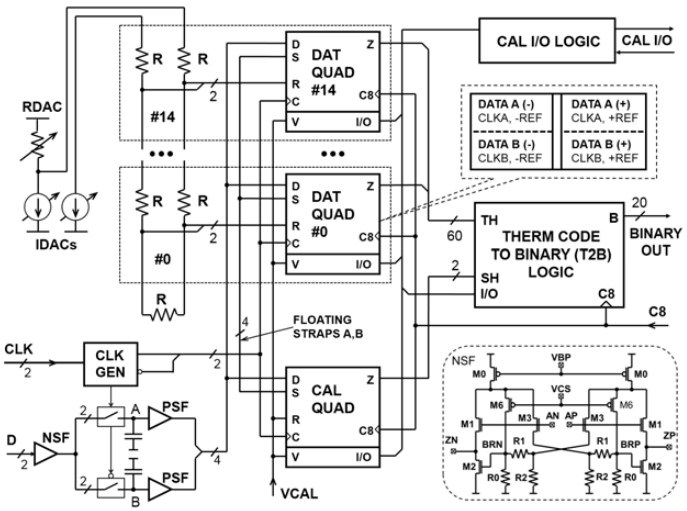

Figure 3.1.2: Architecture of a half-ADC.

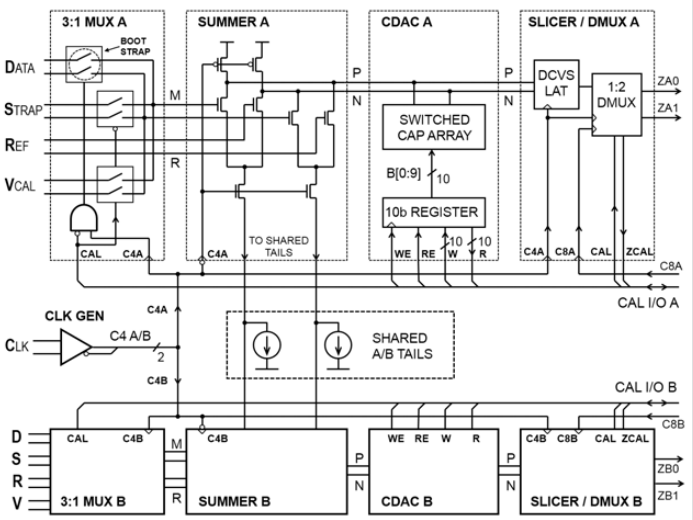

Figure 3.1.3: Block diagram of a comparator pair (A,B).

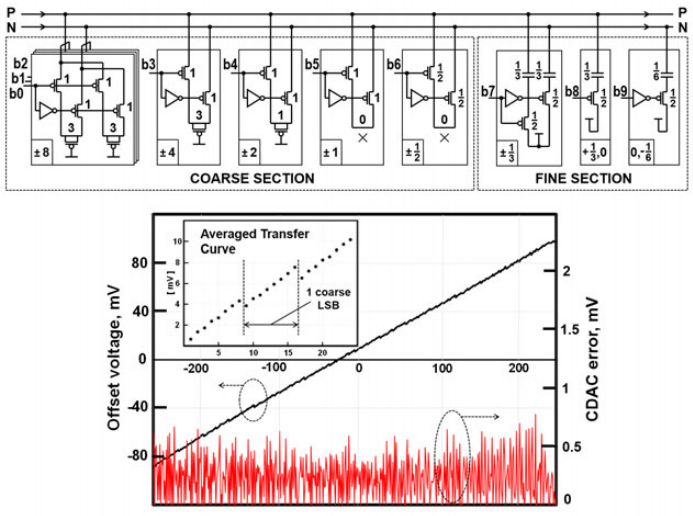

Figure 3.1.4: CDAC schematic and transfer curve.

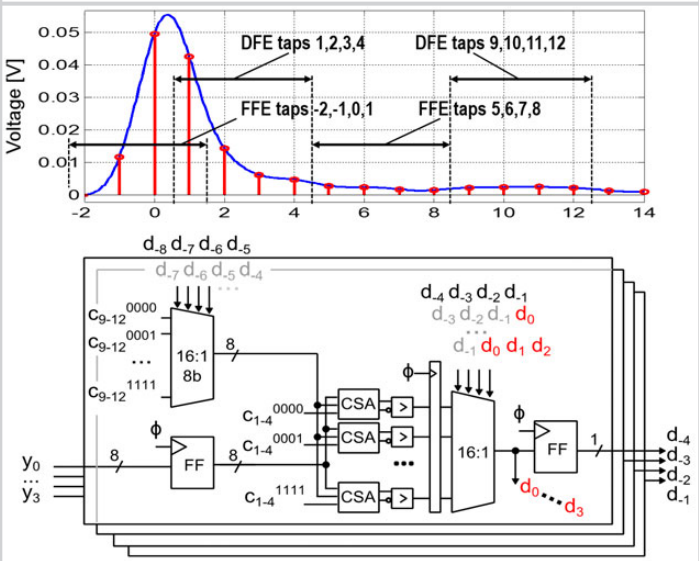

Figure 3.1.5: DFE/FFE partition and DFE block diagram.

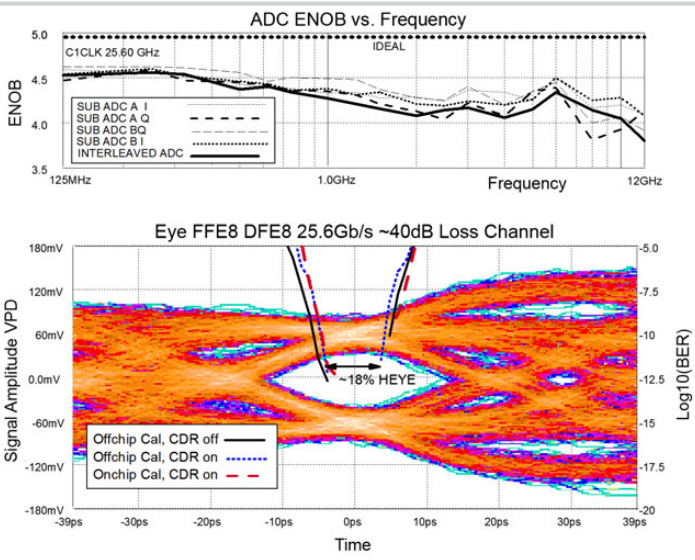

Figure 3.1.6: ADC ENOB, Receiver Eye Diagram and BER.

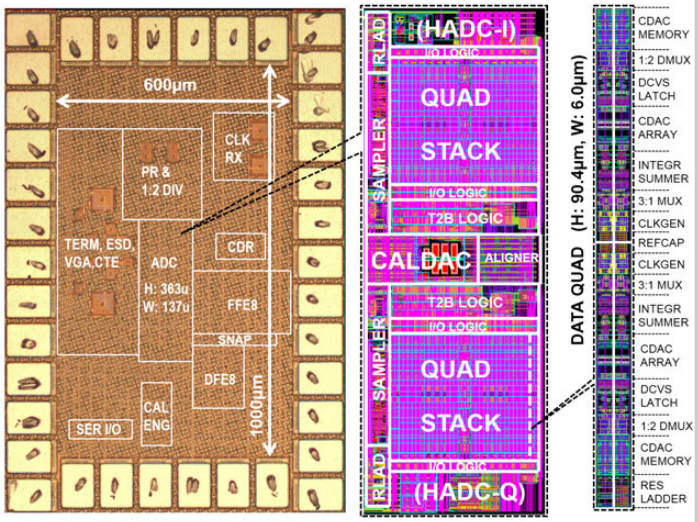

Figure 3.1.7: Die photo with layout detail.

978-1-4673-9467-3/16/$31.00 ©2016 IEEE

3.2 A 320mW 32Gb/s 8b ADC-Based PAM-4 Analog Front-End with Programmable Gain Control and Analog Peaking in 28nm CMOS

Delong Cui[1], Heng Zhang[1], Nick Huang[1,2], Ali Nazemi[1], Burak Catli[1], Hyo Gyuem Rhew[1], Bo Zhang[1], Afshin Momtaz[1], Jun Cao[1]

[1]Broadcom, Irvine, CA,
[2]now with Apple, Cupertino, CA

Multilevel modulation formats, such as PAM-4, have been introduced in recent years for next generation wireline communication systems for more efficient use of the available link bandwidth. High-speed ADCs with digital signal processing (DSP) can provide robust performance for such systems to compensate for the severe channel impairment as the data rate continues to increase.

ADC-based receivers with NRZ format for 10Gbps serial links have been developed [1-3]. For an ADC with a conversion rate of more than 16GS/s, the previous developments have the ADC input directly driving N-way interleaved structures [4-5]. Driving externally allows for higher input amplitude to relax the noise requirement of the ADC. However, the heavy load capacitance presented by the ADC results in degraded return loss and additional insertion loss in the system. Moreover, a programmable gain control (PGA) unit is needed to adapt the input amplitude variation in practical applications. The additional CTLE function can extend the loss compensation capability and reduce the required digital equalizer length and power compared with ADC-only solutions. This paper presents a 32Gbps 8b ADC-based analog front-end with programmable gain control and analog peaking for a PAM-4 wireline system.

The receiver block diagram is shown in Fig. 3.2.1. An input CTLE stage is followed by a two-stage PGA to provide a 14dB gain adjustment and 7dB peaking control. An 8b, four-way time-interleaved ADC digitizes the data for adaptation/equalization and timing recovery. On-chip calibration loops are used to cancel the gain, offset, and timing skew mismatches between the time-interleaved ADC slices.

The circuit diagram of the CTLE and PGA is shown in Fig. 3.2.2. The CTLE stage has a capacitor DAC for peaking control adaptation. The PGA uses resistor degeneration for linearization and the resistor DAC is used for gain control. All the resistor values are calibrated by using an on-chip calibration circuit. Shunt inductors are used to extend the bandwidth while maintaining low power operation. Moreover, to achieve a challenging target of 7b ENOB for the CTLE and PGA, the gain is optimally distributed among the three stages in order to achieve the best trade-off between the kT/C noise and the 11GHz bandwidth requirement.

The 32Gbps receiver has four time-interleaved ADC slices. Each slice is composed of one first level track/hold switch driving eight SAR ADCs as shown in Fig. 3.2.1. After the incoming data is held by the first level switch, the switch at the front of one of the eight 2^{nd}-level SAR ADCs is opened for data to pass through, where it is then held locally for analog-to-digital conversion. During the tracking of this SAR ADC, the seven other SAR ADCs in this slice are in different stages of the SAR conversion cycle.

The 16GHz clock is generated in a PLL and divided down to 8GHz I/Q clocks. After a phase interpolation unit for clock and data recovery, the interleaved clocks are generated and fed into the ADC slices. To minimize the power in the clock path, the clock buffer sizes are reduced. As a result, the random mismatch between the interleaved clocks can be as large as several ps. On the other hand, for a wide-band input with 2^{nd}-order roll off, the required timing mismatch can be calculated as: $\sigma_t^2=M/(M-1)\times 2/3/(2\pi\times f_{3dB})^2/2^{2B}$ [6].

In the above equation, M is the number of interleaved ADCs, f_{3dB} is the bandwidth, and B is the resolution contributed by the skew mismatch with a standard deviation of σ_t. For example, a 0.15ps mismatch in interleaved clocks would degrade the ADC ENOB performance to 7b for an input with a 3dB bandwidth of 8GHz.

As shown in Fig. 3.2.3, a two-step approach is used with a coarse-adjustment capacitor DAC to cover a wide range of random mismatch up to 3ps and a fine-adjustment DAC to provide the required resolution of less than 0.1ps for high-speed operation. Overlap between the coarse adjustment step size and fine adjustment range is guaranteed to ensure the monotonicity and convergence of the search engine. The voltage and temperature drift of the calibration circuit is controlled to be less than 0.1ps. This allows a one-time skew calibration upon the start of the chip and no background calibration is needed to track the environment change to save power.

Unlike standalone ADCs, the total harmonic distortion (THD) of the CTLE and PGA limits the effective input amplitude to the ADC, which in turn limits the signal to noise ratio (SNR) of the whole system. The input amplitude of the SAR unit ADC is chosen to be 400mV for the best tradeoff between the PGA THD and ADC noise floor while maintaining low-power operation.

Each unit SAR ADC uses an asynchronous architecture, as shown in Fig. 3.2.4. The comparator is clocked by a state machine, which starts the next conversion cycle after the previous conversion cycle is finished. The asynchronous structure achieves better power/speed performance compared with its synchronous structure counterparts. A sub-radix, split-capacitor DAC structure is utilized to overcome the error due to capacitor mismatch and reference bouncing. Moreover, by providing over-range protection, incomplete DAC settling is allowed for all conversion cycles except the last two LSBs to maximize the speed of operation while maintaining low power operation for the unit SAR ADC. This, in turn, reduces the number of SAR ADC units to be interleaved, and hence the loading seen by the preceding stages, to reduce the overall system power consumption.

The comparator used in the SAR is based on the strong-arm architecture. Differential reset switches at the drain node of input pairs and comparator outputs increase the reset speed and the comparator speed of operation. The offset calibration is performed by an added differential pair that takes input from a reference ladder, which is controlled by the calibration loop. Compared with the conventional approach of adding programmable caps at comparator outputs, the offset calibration differential pair tracks the input differential pairs as the voltage and temperature drift. As a result, this topology offers a better voltage and temperature tracking capability, which can be in the range of ±0.5 LSB.

The chip is fabricated in 28nm CMOS. The measured ENOB of the receiver is 6.4b at DC and 5.85b at Nyquist, as shown in Fig. 3.2.5. The receiver SNDR performance includes the contribution from both the CTLE/PGA and ADC, which contributes equally at DC. This results in an ADC-only ENOB of ~7b at DC. The receiver consumes 320mW of power and achieves a figure of merit (FOM = power/($f_s\times 2^{ENOB}$(Nyquist)) of 0.35pJ/conversion-step. The receiver has been tested to achieve error-free operation with various backplanes and cables and other low-cost medium channels, such as Thunderbolt, eSATA, CX1, and USB, with loss of up to 32dB at 8GHz and more than 50dB at 16GHz. Figure 3.2.6 compares this ADC-based PAM-4 analog front-end with other published results with similar data rates. Measured channel compensation results are summarized for other NRZ/PAM-4 ADC-based analog front ends and this PAM-4 based design. The receiver area is 0.89mm^2.

References:

[1] M. Harwood et al., "A 12.5Gb/s SerDes in 65nm CMOS Using a Baud-Rate ADC with Digital Receiver Equalization and Clock Recovery", *ISSCC Dig. Tech. Papers*, pp. 436-437, Feb. 2007.
[2] S. Verma et al., "A 10.3GS/s 6b flash ADC for 10G Ethernet applications", *ISSCC Dig. Tech. Papers*, pp. 462-463, Feb. 2013.
[3] B. Zhang et al., "A 195mW / 55mW dual-path receiver AFE for multi-standard 8.5-to-11.5 Gb/s serial links in 40nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 544-545, Feb. 2013.
[4] C.-C. Huang et al., "A CMOS 6-Bit 16-GS/s Time-Interleaved ADC with Digital Background Calibration", *IEEE Symposium on VLSI Circuits Dig. Tech. Papers,* pp. 159-160, June 2010.
[5] L. Kull, et al., "A 90GS/s 8b 667mW 64× Interleaved SAR ADC in 32nm Digital SOI CMOS", *ISSCC Dig. Tech. Papers*, pp. 378-379, Feb. 2014.
[6] M. El-Chammas, et al., "General Analysis on the Impact of Phase-Skew in Time-Interleaved ADCs", *IEEE Transactions on Circuits and Systems-I: Regular Papers,* vol. 56, no. 5, pp 902-910, May 2009.

 978-1-4673-9467-3/16/$31.00

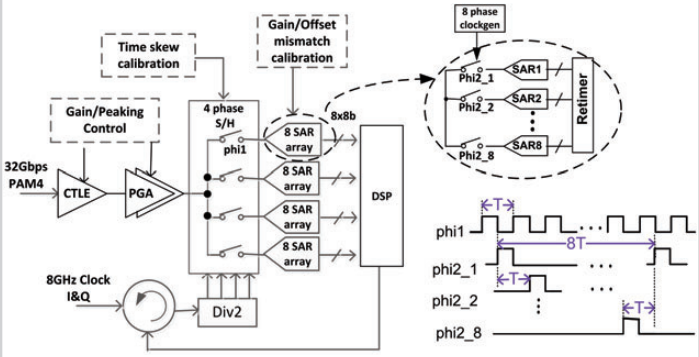

Figure 3.2.1: 16GS/s ADC-based receiver architecture.

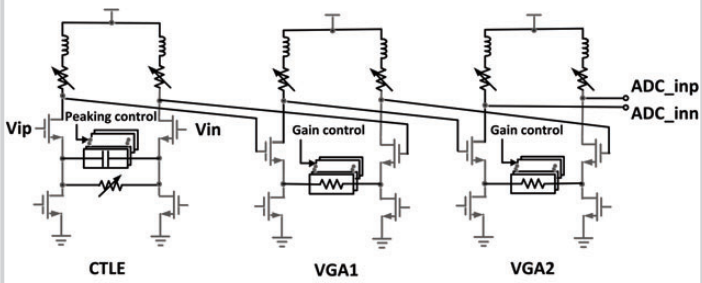

Figure 3.2.2: Circuit diagram of the CTLE/PGA.

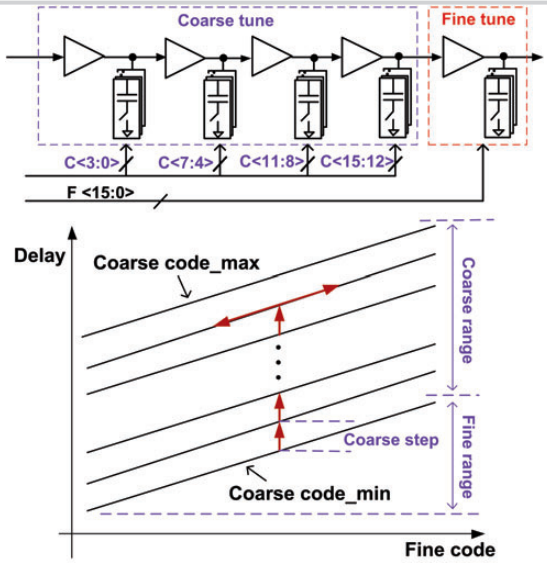

Figure 3.2.3: Timing skew adjustment circuit.

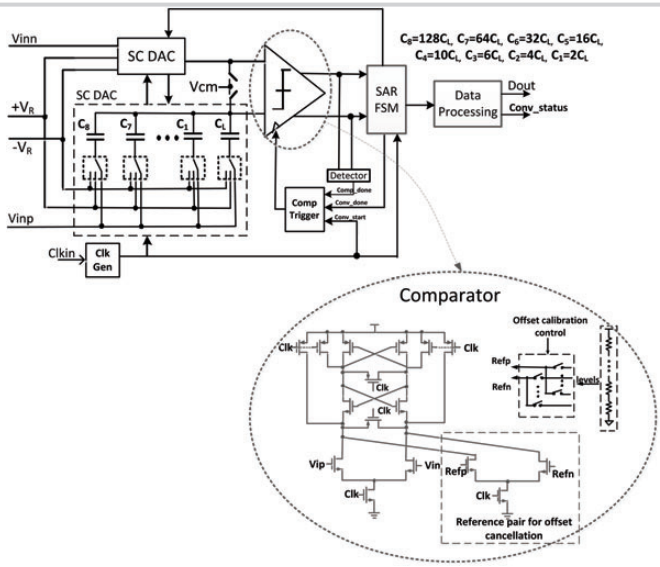

Figure 3.2.4: Unit asynchronous SAR ADC architecture.

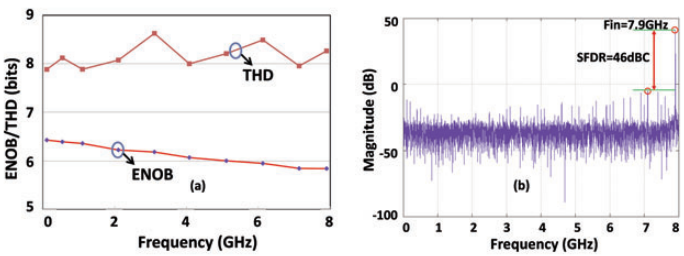

Figure 3.2.5: (a) ENOB performance versus frequency, (b) spectrum at Fin = 7.9GHz.

Design	[1]	[2]	[3]	[4]	This work
Technology	65nm	40nm	40nm	65nm	28nm
Sampling Rate (GS/s)	12.5	10	10	16	16
Data Format	NRZ	NRZ	NRZ	N/A*	PAM4
Data rate (Gbps)	12.5	10	10	N/A	32
Resolution (bit)	4.5	6	6	6	8
Power (mW)	230	240	195	435*	320
ENOB @ DC	N/A	5.5	4.9	4.35	6.4
ENOB @ Nyquist	N/A	5.1	4.6	N/A	5.85
Compensated Channel Loss at Nyquist	15dB	35dB	34dB	N/A*	>50dB
FOM (fJ/c-s)	N/A	700	832	2600	350

* ADC Only, No VGA/CTLE

Figure 3.2.6: Performance of >10Gbps ADC-based receivers.

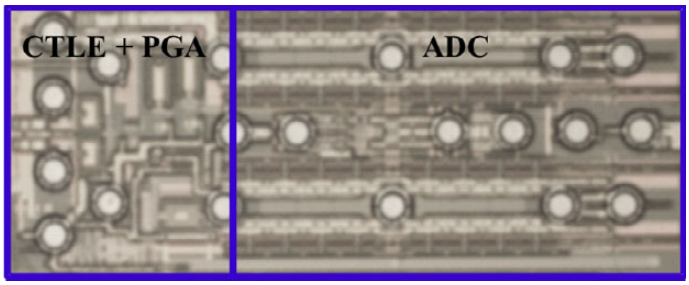

Figure 3.2.7: Die micrograph.

978-1-4673-9467-3/16/$31.00

3.3 A 25Gb/s Multistandard Serial Link Transceiver for 50dB-Loss Copper Cable in 28nm CMOS

Takayasu Norimatsu[1], Takashi Kawamoto[1], Kenji Kogo[1], Naohiro Kohmu[1], Fumio Yuki[1], Norio Nakajima[2], Takashi Muto[2], Junya Nasu[2], Takemasa Komori[2], Hideki Koba[2], Tatsunori Usugi[2], Tomofumi Hokari[2], Tsuneo Kawamata[2], Yuichi Ito[2], Seiichi Umai[2], Masatoshi Tsuge[2], Takeo Yamashita[2], Masatoshi Hasegawa[2], Keiichi Higeta[2]

[1]Hitachi, Tokyo, Japan, [2]Hitachi, Kanagawa, Japan

The amount of data traffic is increasing year by year as the number of data-rich services like cloud services and streaming services are increasing. The number of switch modules between servers should decrease to lower latency, and several servers in each rack should be connected to one switch module with cables in a data centre. Using copper cables to connect racks is attractive in terms of cost minimization. Thin cables, for example 34 AWG copper cables, make maintenance easy. The cable length should be 5–7m to connect between racks, and 34 AWG 7m cable has 48dB loss, including board trace loss, package loss and so on. So far transceivers over 25Gb/s, equalizing 35–40dB channel loss have been proposed [1-4], with which low-loss cables like 26 AWG have been required. We target a 25Gb/s transceiver equalizing over 50dB channel loss, and adopt a sub-mV dynamic DC offset cancelation and a decision-feedback equalizer (DFE) with a bias-controlled tap slicer. Both improve on the minimum input sensitivity and enable data transmission through a channel with over 50dB loss.

The transceiver architecture is shown in Fig. 3.3.1. This transceiver has 8-lane TX and RX blocks and a logic block. Each lane has three equalizers: a 4-tap FFE in the TX block, and a CTLE and 14-tap DFE in the RX block. The tap coefficients of the DFE are calculated automatically with a tap calculation block (TAP calc.). The phase of clocks for the DFE is adjusted adaptively with clock and data recovery (CDR) logic and a phase interpolator (PI). The clocks are distributed from an RXPLL to the PI through a duty-cycle corrector (DCC) in the RX block. The divided clock is distributed from the PI to a TXPLL as a reference clock to synthesize the clock phase with the phase of the input data in the RX block. Both TXPLL and RXPLL have a ring oscillator for less than 7GHz and a LC oscillator for more than 12GHz. To reduce power consumption, RXPLL and TXPLL are shared by 4-lane TX and RX blocks.

A CTLE is the most effective block to improve SNR in the RX. High gain at the Nyquist frequency, low input referred noise, and low DC offset are required for a CTLE. The CTLE block diagram is shown in Fig. 3.3.2. The CTLE gain at 12.78GHz is designed to be over 20dB to have enough SNR at the DFE input, so this CTLE has 6-stage amplifiers. The first amplifier in the CTLE has enough gain to lower the NF of the CTLE. The CTLE has 2-zero peak amplifiers, which have a peak at low frequency, for example 1GHz, and lower jitter without requiring a DFE with many taps [5]. The peak amplifier has programmable resistors, a capacitor, and a series of programmable resistors and a capacitor in parallel at the source of the MOSFETs as shown in Fig. 3.3.2, which makes 2 zeros and 1 pole. The signal level through the 50dB-loss channel is very low; the $1.1V_{p\text{-}p}$ differential TX output is attenuated to a $3.5mV_{p\text{-}p}$ differential RX input. Thus, the DC offset is targeted to be less than 1mV at the RX input. DC offset is compensated at 4 points as shown in Fig. 3.3.2 to suppress DC offset variation due to gain change of the peak amplifiers and VGAs. The DC offset is coarsely compensated with this compensation, and the residual DC offset is less than 3mV. In addition to this, a dynamic offset cancelation feedback loop is adopted to suppress DC offset variation due to environmental conditions like supply voltage and temperature. As shown in Fig. 3.3.2, the comparator in the dynamic DC offset cancelation block does not add DC offset. The DC offset remaining after trimming is reduced to less than 1mV with the dynamic DC offset cancelation feedback loop.

Figure 3.3.3 illustrates the 14-tap DFE block diagram. The 14-tap DFE has EVEN and ODD blocks and each of them operates at half rate. The former 2 taps are unrolled, which is a popular architecture for high-speed DFE [1-4], and taps 3 to 14 are added with a current summer. Tap 3 is added to the outputs of the 1st common mode logic (CML) directly because it has little timing margin. Taps 4 to 14 are added to outputs through cascode MOSFETs to decrease the load capacitance at outputs of the 1st CML [6]. The currents of taps 4 to 14 are added at a cascode MOSFET source to decrease switching delay by adding currents directly to the outputs of the 1st CML.

To judge whether the attenuated input signal is high or low accurately, slicers in a DFE should have high sensitivity. Especially if part of the slicers are used as a phase detector for CDR and are required to detect data edges. Low delay is also required for slicers. For example, the delay from the slicer to slicer for tap 3 should be less than 2.5UI (89ps in 28Gbps). To satisfy these requirements, we propose a bias-controlled tap slicer that controls the tap and DC offset compensation with bias current, as shown in Fig. 3.3.4. This architecture removes the current summer from slicer outputs in contrast to previous slicers [1,3-5], and so lowers the delay and minimum input sensitivity due to smaller parasitic capacitance. Taps 1 and 2 and the DC offset compensation value are added to the slicer output as bias current. The slicer adopts dynamic logic architecture for low delay, which charges while CLK is low and compare differential inputs while CLK is high. The minimum input voltage of the slicer is less than 1.2mV with less than 30ps delay in simulation, which is equal to a 1.6ps dead zone of CDR if the total gain at the Nyquist frequency, including channel loss and CTLE gain, is -30dB, and the amplitude of TX differential outputs is $1.1V_{p\text{-}p}$. In addition to this, the bias current of the slicer always flows for low delay while it is turned on. This can reduce current consumption because we can remove large clock buffers to drive the bias current of the slicers.

The transceiver is fabricated in a 28nm CMOS process and the performance of the transceiver is measured in the test setup shown in Fig. 3.3.5. 5m 34 AWG QSFP copper cable is used to communicate between transceivers. The channel characteristic of BGA to BGA in this test setup is shown in Fig. 3.3.5. The channel loss at 12.89GHz is about 51dB including the 15dB evaluation board trace loss. The PRBS31 data pattern at 25.78Gb/s is transmitted in this test setup, and an internal eye monitor output in Fig. 3.3.6 shows that eye height is over 57mV and eye width is over 0.3UI. The adopted high sensitivity tap bias control slicer and dynamic DC offset cancelation make the minimum input sensitivity less than 3mV when 25.78Gb/s 1/0 data pattern is input. Therefore, as shown in Fig. 3.3.6, the BER is less than 1E-12 when the dynamic DC offset cancelation block is enabled, whereas the BER is over 1E-10 when the dynamic DC offset cancelation block is disabled. The transceiver consumes 403mW from a 0.9V and a 1.5V supply per lane, which includes 1/4 PLL power consumption and 1/8 common block power consumption. This chip has 8-lane transceivers and the chip area is 5.5×4.6 mm^2 as shown in Fig. 3.3.7.

References:

[1] J. Bulzacchelli et al., "A 28Gb/s 4-Tap FFE/15-Tap DFE Serial Link Transceiver in 32nm SOI CMOS Technology", *ISSCC Dig. Tech. Papers*, pp. 324-325, Feb. 2012.

[2] H. Kimura et al., "28Gb/s 560mW Multi-Standard SerDes with Single-Stage Analog Front-End and 14-Tap Decision Feedback Equalizer in 28nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 38-39, Feb. 2014.

[3] T. Kawamoto et al., "Multi-Standard 185fsrms 0.3-to-28Gb/s 40dB Backplane Signal Conditioner with Adaptive Pattern-Match 36-Tap DFE and Data-Rate-Adjustment PLL in 28nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 54-55, Feb. 2015.

[4] B. Zhang et al., "A 28Gb/s Multi-Standard Serial-Link Transceiver for Backplane Application in 28nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 52-53, Feb. 2015.

[5] S. Parikh et al., "A 32Gb/s Wireline Receiver with a Low-Frequency Equalizer, CTLE and 2-Tap DFE in 28nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 28-29, Feb. 2013.

[6] C. Thakkar et al., "A Mixed-Signal 32-Coefficient RX-FFE 100-Coefficient DFE for an 8Gb/s 60GHz Receiver in 65nm LP CMOS", *ISSCC Dig. Tech. Papers*, pp. 238-239, Feb. 2013.

 978-1-4673-9467-3/16/$31.00

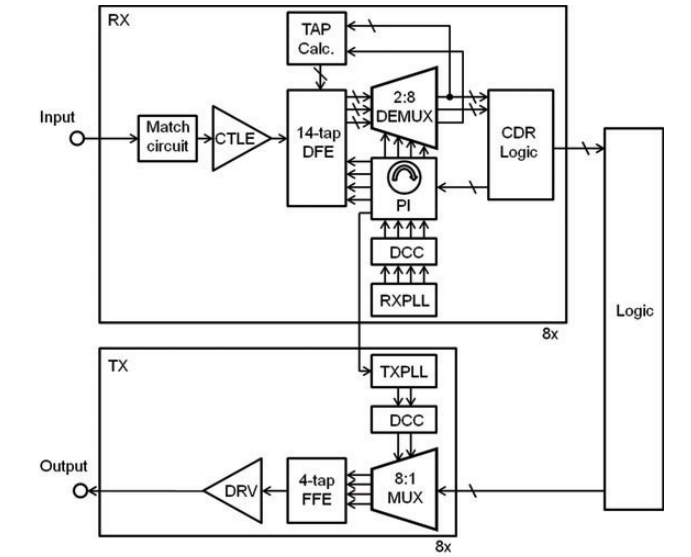

Figure 3.3.1: Transceiver block diagram.

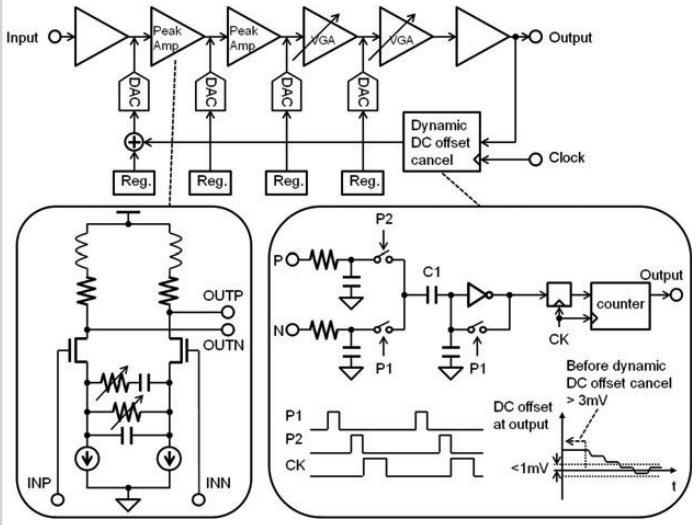

Figure 3.3.2: CTLE block diagram with dynamic DC offset cancelation.

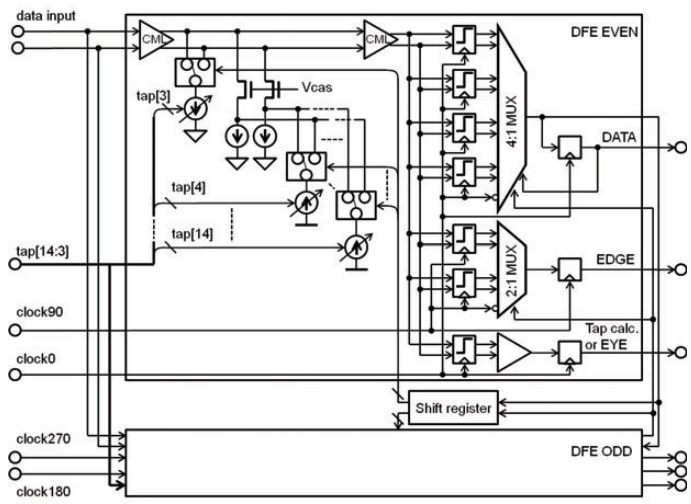

Figure 3.3.3: 14-tap DFE block diagram.

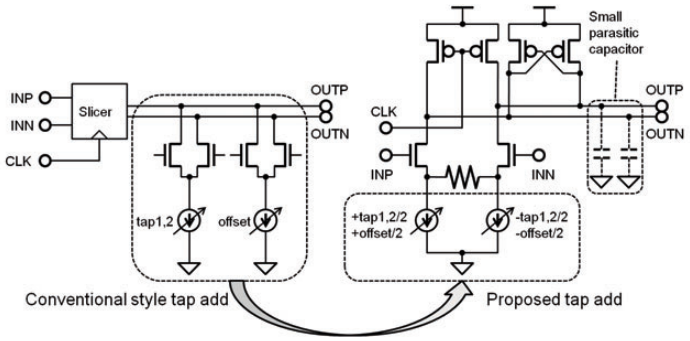

Figure 3.3.4: Slicer block diagram.

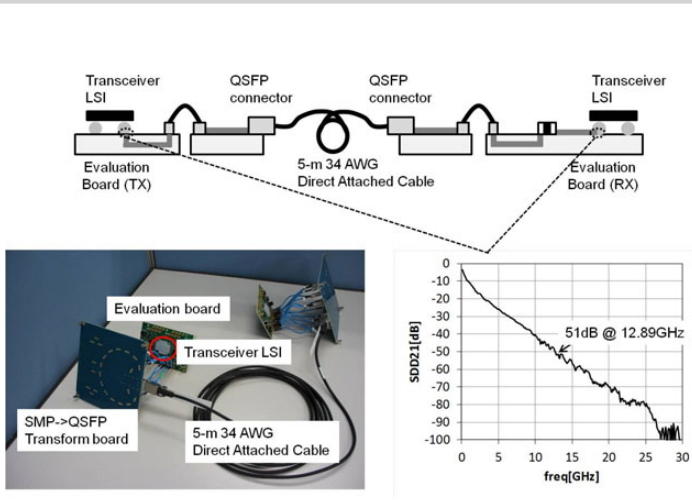

Figure 3.3.5: Evaluation setup and channel characteristic.

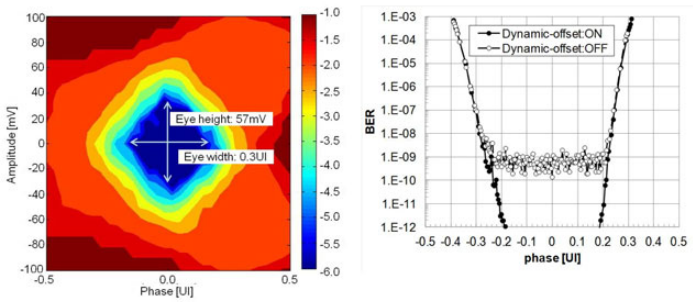

Technology	28-nm CMOS
Supply voltage	0.9 / 1.5 V
Loss at Nyquist	51 dB (@ 12.89GHz)
Data rate	0.3 – 28 Gbps
Power consumption per lane	403 mW
Chip area	5.5 x 4.6 mm^2

Figure 3.3.6: Measurement results.

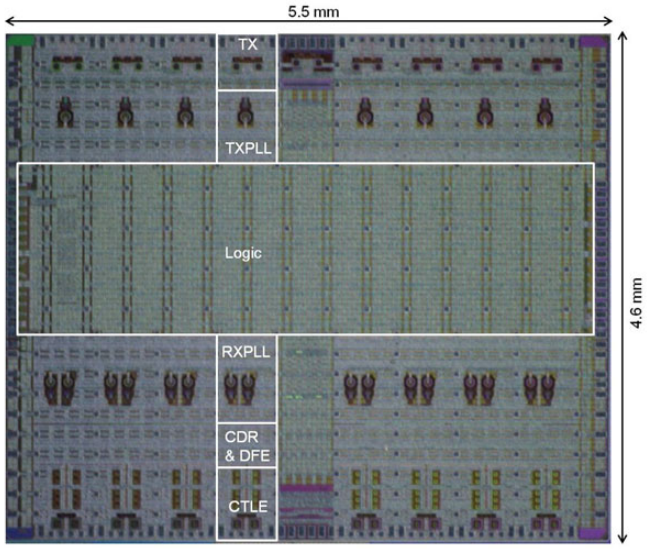

Figure 3.3.7: Chip micrograph.

978-1-4673-9467-3/16/$31.00

3.4 A 40/50/100Gb/s PAM-4 Ethernet Transceiver in 28nm CMOS

Karthik Gopalakrishnan[1], Alan Ren[1], Amber Tan[1], Arash Farhood[1], Arun Tiruvur[1], Belal Helal[1], Chang-Feng Loi[2], Chris Jiang[1], Halil Cirit[1], Irene Quek[2], Jamal Riani[1], James Gorecki[1], Jennifer Wu[1], Jorge Pernillo[1], Lawrence Tse[1], Michael Le[3], Mohammad Ranjbar[1], Pui-Shan Wong[1], Pulkit Khandelwal[1], Rajesh Narayanan[1], Ravindran Mohanavelu[1], Sameer Herlekar[1], Sudeep Bhoja[1], Vlad Shvydun[1]

[1]Inphi, Santa Clara, CA,
[2]Inphi, Singapore, Singapore, [3]Inphi, Irvine, CA

High-speed signaling using NRZ has approached speeds above 50Gb/s where it is extremely difficult to maintain power efficiency and performance over a wide variety of channels and applications. PAM-4 is emerging as one way to increase throughput in such band-limited channels. Higher modulation formats help to address cost in optical systems by packing more bits/wavelength [1]. Strong momentum in standards to adopt PAM-4 reflects these significant trends in the industry. At the same time, migrating transceiver designs to current technology nodes have narrowed the power gap between traditional Analog and ADC-DSP-DAC-based systems for high-speed applications. These factors make ADC-based receivers a highly desirable choice, as is also the trend in wireless communications.

The transceiver shown in Fig. 3.4.1 is built to support a wide range of Ethernet applications. The block diagram of the complete receiver is shown in Fig. 3.4.2. A merged equalizer-VGA (Fig. 3.4.2b) performs coarse equalization and drives the T/H circuits, minimizing the number of stages. The VGA has a gain range >12dB in 0.1dB steps and a BW >13GHz. A high-frequency gain-boost of up to 8dB provides system SNR improvement in high loss channels. This equalizer-VGA achieves SNR >41.7dB and THD >36dB over all gain, boost and process corners by employing g_m-boosted source degeneration [2]. G_m of transistors P1 and P2 is boosted by the loops created by N1, N3, N7 and N2, N4, N8, respectively. Programmable gain is achieved through differentially modulating the mirrored transconductance gain by creating a differential V_{ds} bias on N3, N5 and N4, N6. Each VGA drives 4 T/H switches operating at 3.5GS/s each. A unity-gain buffer is used at the output of each track-and-hold to drive four sample-and-hold circuits and their associated capacitive DACs. A feed-forward-based negative-g_{ds} technique in an otherwise conventional source follower helps to optimize signal-to-noise performance and linearity at low supply voltages. A replica circuit controls the negative-g_{ds} in the buffers ensuring constant gain over process, voltage, and temperature. Each of the 32 sub-ADCs is a SAR core clocked at 7GHz. Independent reference buffers minimize non-linear and signal dependent noise coupling between channels.

The DSP core (Fig. 3.4.3) performs the calibration for offset, gain, timing skew estimation and correction of the analog front-end. The DSP core employs a set of parallel FFEs for channel equalization. The parallel factor was chosen to be a multiple of the number of sub-ADC channels to optimize power consumption. The gain of the 32 time-interleaved channels is estimated using an envelope detector. Any gain mismatch is compensated by adjusting the associated reference voltage, which maximizes range of each ADC slice. Residual gain errors are further corrected in the DSP. Offsets of each interleaved channel are estimated digitally by computing the average of the slicer error at the FFE output that corresponds to each signal path and corrected inside the DSP core. Here, we trade-off dynamic range of the ADC for offset correction to avoid DACs in the signal path that degrade bandwidth. Timing mismatch is estimated digitally by using correlated properties of the PAM input signal. Digital controls are fed back to small delay cells that alter the sampling phase of the 8 T/H clocks with a resolution of ~100fs. Figure 3.4.3 plots SNDR of the entire receive chain and the impact from timing calibration. BW mismatch between the different AFE paths is compensated by independent adaptation of the FFE slices.

The DSP employs an adaptive PAM-4 DFE. The feedback taps are limited to one tap to minimize impact of error propagation. Baud rate clock recovery is used which is based on the well-known Mueller-Muller timing recovery scheme [3] but taking inputs directly at the ADC output, thus eliminating interaction problems with FFE-DFE adaptation while providing a low latency clock recovery path. A measured jitter tolerance plot for NRZ modulation is shown in Fig. 3.4.3 against a VSR mask. The clock recovery scheme can be made truly reference-less by taking advantage of the reference-less HOST VSR link. The recovered clock is filtered prior to ADC sampling. This allows the reference to be powered down.

A DLL, as shown in Fig 3.4.4a, generates the timing phases for the ADC from a 7GHz clock from the PLL. Static phase offset (SPO) is a well-known problem in DLLs. A differential quadrature phase detector achieves the goal of low SPO, allowing for healthy timing margins in the ADC clocking and easing start up of the DSP engine.

The line transmitters can be configured as two PAM or four NRZ links. The latter provides support for segmented modulators that generate PAM-4 in the optical domain. The transmitter implementation, shown in Fig. 3.4.5a, is a traditional CML implementation with shunt peaking in the final stage. The driver provides swing levels up to 1.4Vpp and incorporates a 3-Tap FIR filter with independent control on the MSB and LSB paths. The MSB to LSB ratio can also be altered to provide pre-distortion on the PAM-4 eye, which is useful in applications where the PAM transmitter interfaces with optical drivers. Figure 3.4.5 shows the transmit eye diagrams measured at the chip balls.

A skew control mechanism shown in Fig. 3.4.5b, auto-zeroes electrical and logical skew in NRZ mode. It can also pre-compensate skews (< ±1UI) that occur downstream. An analog phase detector senses the skew and a FSM corrects for it by adjusting the offset in the PLL charge-pump. A delta-sigma modulator driving this offset current provides very fine control of the PLL phase (resolution <100fs). The FSM also calibrates the loop to be able to introduce any required skew and maintain it over operating conditions. The noise introduced by the offset delta sigma is negligible. The entire system shows a simulated 3σ error of < +/- 0.5ps peak-to-peak due to random mismatches.

A fractional-N PLL provides the required clocks for the TX and RX paths and is shown in Fig 3.4.6. The VCO is inductor-based with dual tuning paths (9.9 to 15.5GHz). A simple amplifier and RC filter form a slow path that drives Vctrl_fast close to a target voltage. This has several advantages: maximizing charge pump headroom and linearity, stabilizing the fast loop K_{vco} over tuning range, tracking temperature and reducing the loop filter size. The multi-modulus (MM) divider is based on Vaucher's extended range topology. This implementation enables transition across stage boundaries, smoothly overcoming a key limitation in the original topology. Figure 3.4.6 shows integrated rms jitter of 181fs on the TX outputs in a frequency band from 1kHz to 100MHz.

Power supply noise management is a key aspect of high performance communication links. Both PSRR and random noise from regulators impact overall SNR of the analog front-ends. The regulator topology shown in Fig. 3.4.4b uses feed-forward injection. The frequency of injection is tuned to attenuate external switching regulator noise, which can often occur around PLL corner frequencies. This attenuation allows for reduced on-board filtering requirements. Source degeneration is employed in the error amplifier to further reduce 1/f noise contributors. The chip is fabricated in 28nm CMOS. The chip consumes 2.4W with 100G traffic and 25Gbaud PAM-4 on the line while the DSP was set to 10-tap FFE configuration.

Acknowledgements:
We thank Aaron Buchwald for providing insightful suggestions. We also thank George Thomas, Goutham Mallareddy and Ishwar Hosagrahar for all lab measurements.

References:
[1] M. Traverso et al., "25GBaud PAM4 Error free transmission over both SMF and MMF in a QSFP form factor on Silicon Photonics," *Optical Fiber Communication Conference*, Th5 B.3, 2015.
[2] S. D. Willingham and K. Martin, "Integrated Video-Frequency Continuous-Time Filters: High-Performance Realizations in BiCMOS," Springer, 1995.
[3] K. H. Mueller and M. S. Muller, "Timing Recovery in Digital Synchronous Data Receivers," *IEEE Trans. on Communications*, vol. COM-24, pp. 516-531, May 1976.

 978-1-4673-9467-3/16/$31.00

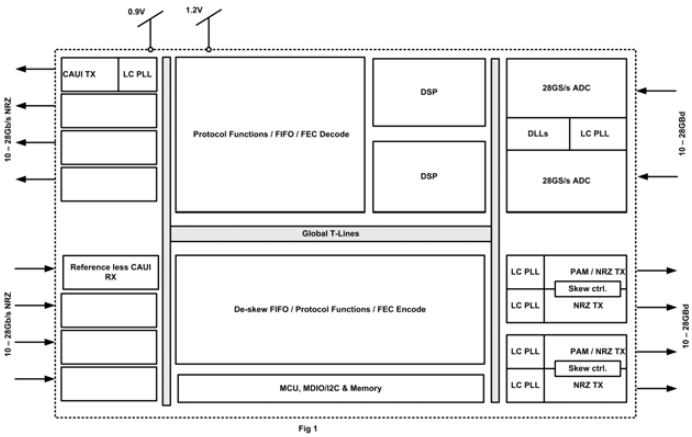

Figure 3.4.1: Transceiver Block diagram.

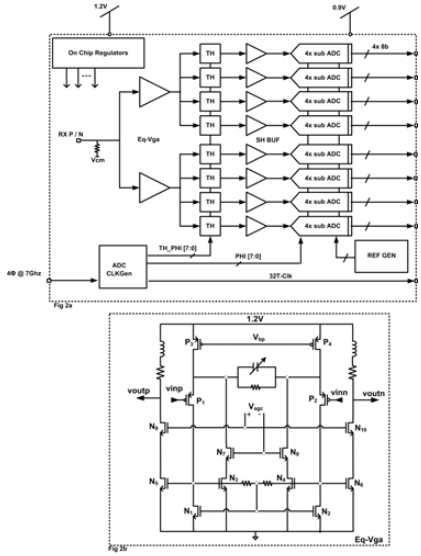

Figure 3.4.2: RX and equalizer-Vga.

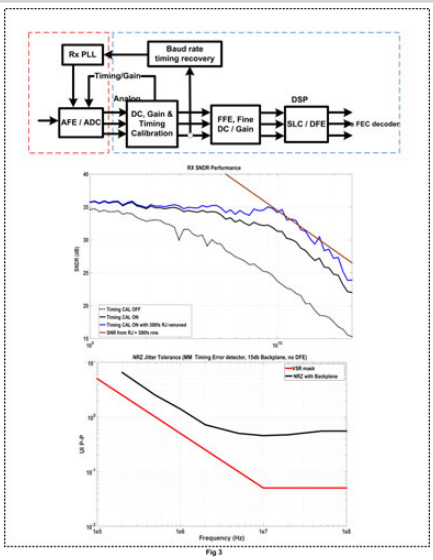

Figure 3.4.3: DSP Block diagram.

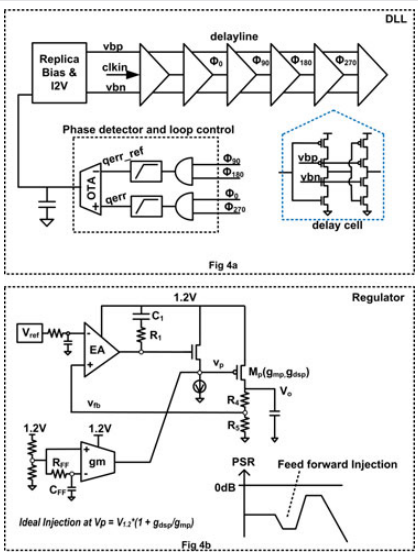

Figure 3.4.4: DLL and Regulator.

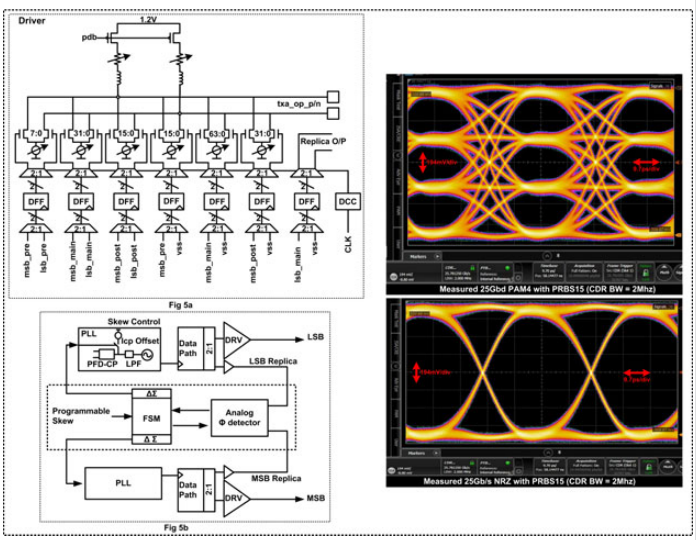

Figure 3.4.5: TX Block diagram.

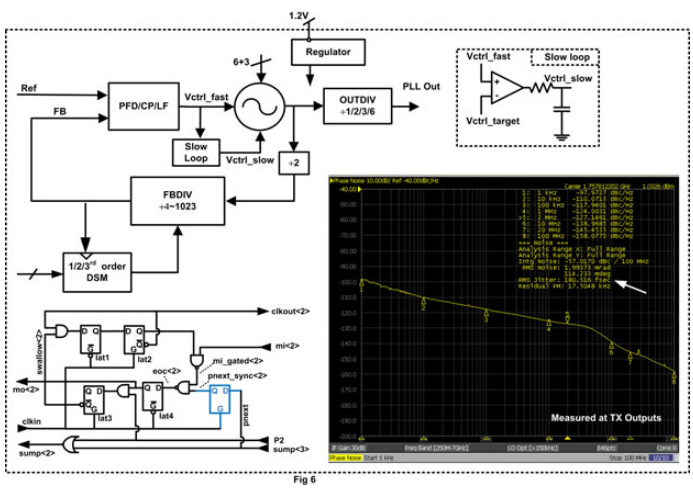

Figure 3.4.6: PLL block diagram.

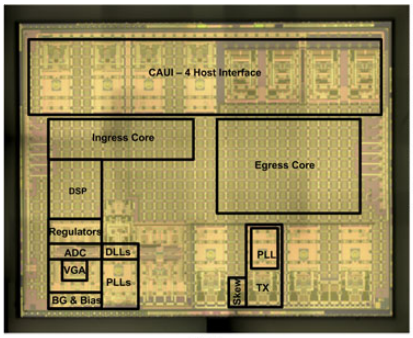

Fig 7

Performance Summary		
Max Data rate	32	Gbd
RX Analog BW	18	Ghz (Typical)
AFE ENOB	5.7	@ 1Ghz
	4.9	@ Nyquist
Jitter Tolerance	Exceeds VSR	Measured for NRZ
PPM Tracking	< +/- 400	ppm w.r.t Reference Clk
Max TX Swing	1400	mVppd
NRZ Skew Control	+/- 1	UI
Skew Resolution	100	fs
TX PLL RJ	181	fs rms (1 kHz - 100 MHz)
TX RJ	240	fs rms (> 2 MHz)
Supplies	1.2 / 0.9	Volts
Power	2.4	Watts with 100G Ethernet Traffic
Chip Area	6.3 x 4.9	mm
Technology	28nm	CMOS Logic

Figure 3.4.7: Die Photo and Performance summary.

978-1-4673-9467-3/16/$31.00

3.5 A 56Gb/s NRZ-Electrical 247mW/lane Serial-Link Transceiver in 28nm CMOS

Takayuki Shibasaki[1], Takumi Danjo[1], Yuuki Ogata[1], Yasufumi Sakai[1], Hiroki Miyaoka[2], Futoshi Terasawa[2], Masahiro Kudo[2], Hideki Kano[2], Atsushi Matsuda[2], Shigeaki Kawai[2], Tomoyuki Arai[2], Hirohito Higashi[2], Naoaki Naka[2], Hisakatsu Yamaguchi[1], Toshihiko Mori[1], Yoichi Koyanagi[1], Hirotaka Tamura[1]

[1]Fujitsu Laboratories, Kawasaki, Japan,
[2]Socionext, Yokohama, Japan

With the rapid growth of data traffic in data centers, data rates over 50Gb/s/signal (e.g., OIF-CEI-56G-VSR) will eventually be required in wireline chip-to-module or chip-to-chip communications [1–3]. To achieve better power efficiency than that of existing 25Gb/s/signal designs, a high-speed yet energy-efficient front-end is needed in both the transmitter and receiver. A receiver front-end with baud-rate architecture [1] has been successfully operated at 56Gb/s, but additional components such as eye-monitoring comparators, phase detectors, and clock recovery circuitry as well as a power-efficient transmitter are needed to build a complete transceiver.

The transceiver presented in this work has a two-lane configuration, with each lane supporting a 56Gb/s non-return-to-zero (NRZ) signal transfer. Two transmitter lanes share a half-rate clock-generation unit (CGU) embedded in between the lanes and each receiver lane has a dedicated quarter-rate clock-recovery unit (CRU) (Fig. 3.5.1). The CGU generates a differential 28GHz clock signal for the 2-to-1 multiplexers (MUXs) in the transmitter front-end by means of a VCO-based PLL. It also generates four-phase 14GHz clock signals that drive the 32-to-2 MUXs preceding the 2-to-1 MUXs by dividing the 28GHz clock signal. The CRU produces a four-phase 14GHz clock signal for the quarter-rate front-end of the receiver by dividing the 28GHz clock signal supplied from a VCO. Both the CGU and CRU are arranged so that an inductor-loaded narrow-band buffer sends the clock signal directly to the target blocks (Fig. 3.5.2).

The incoming receiver signal first goes through a T-coil-based termination block and enters into a continuous-time linear equalizer (CTLE). The CTLE consists of a high-frequency equalizer (HFEQ) using a single-stage source-degenerated amplifier, a two-stage low-frequency equalizer (LFEQ) in which a feedback circuit generates an additional zero at a low frequency, and a single-stage output [4]. The amplifier stages in the forward path have a shunt-peaking inductor to enhance the bandwidth. The –3dB bandwidth of the CTLE is over 30GHz with a DC gain of 14dB. The CTLE output is digitized by 4-way-interleaved high- and low-threshold comparator arrays (i.e., eight comparators altogether) driven by the four-phase 14GHz clock signals, performing baud-rate sampling. The outputs of the comparator arrays are used for both speculative decision-feedback equalizing and the phase detection, requiring two comparators activated per unit interval (UI) [1]. An additional comparator driven by the 14GHz phase-adjustable clock signal supplied from a phase interpolator (PI) is used for the equalizer adaptation and eye monitoring.

The outputs of the nine (8+1) comparators are first demultiplexed by an 8-to-16 and a 1-to-2 demultiplexer (DMX), resulting in an 18b, 7Gb/s signal. The 16b portion of the signal, which originated from the high- and low-threshold comparator arrays, is sent to the phase detector (PD) that detects the phase difference between the input data and recovered clock signal.

A 1-tap decision-feedback equalizer (DFE) selector adopting a parallelized look-ahead technique [1] selects a decision bit from the two 16b words using two stages of selector arrays. The first stage generates two decision-bit sequences assuming that the preceding bit in each sequence is correct. The correct decision for the first bit of the sequence is known from the last decision bit of the previous cycle and the second stage selects the correct sequence out of the two accordingly.

The PD operates in a 7GHz clock domain and generates a two-bit output, PDOUT, which indicates whether the clock phase should go up, go down, or stay put. Eight PD-logic blocks output the phase difference when the three-bit decision sequences '011' or '100' are detected (Fig. 3.5.3). For the PD to operate correctly, the input signal should be equalized such that the cut-off frequency of the signal is about one quarter to one third of the baud frequency. We confirmed by simulations that the PD detects the center phase of the decision input eye with a phase error less than 0.1UI when the cut-off frequency of the signal at the CTLE output is between 14 and 18GHz. The PD-logic block generates two-bit UP/DN signals that use a redundant binary representation where "10" represents "up", "00" represents "down", and "01" represents "stay". The UP/DN signals from the PD-logic blocks are summed and input to a quantizer. Since "up" and "down" are always followed by "stay" and the redundant binary representation is used, the gate count of the adder is reduced, enabling 7GHz operation. The quantizer sums the input with the residual quantization error and stores the result in a register. The upper two bits of the register represent the quantized output and the lower two bits represent the residual error. Since the quantization error is not discarded by a truncation, the PD signal-to-noise ratio is not sacrificed and the charge-pump design is simplified.

The transmitter generates a 56Gb/s signal by using 2-to-1 MUXs driven by the half-rate clock signal supplied from the CGU. A 32-to-2 MUX preceding the 2-to-1 MUXs outputs two pairs of 28Gb/s data needed to operate the 2-tap feed-forward equalizer (FFE). The 2-to-1 MUX consists of two tri-state buffers with the output nodes tied together. A tri-state buffer with a NAND/NOR input stage is used (Fig. 3.5.4). The driver stage is a source-series-terminated (SST) driver that has a capacitor connected in parallel with the series-termination resistor to enhance the eye opening [5]. The 2-tap FFE is implemented by connecting the outputs of two SST drivers composed of 64 identical slices for main-tap and 32 identical slices for pre-tap. The number of active slices in the SST driver is adjusted depending on the FFE strength and termination resistor.

To minimize the transmitter power consumption, reducing the capacitive load for the 28GHz clock signal is crucial, since the 28GHz clock signals are of the highest frequency in the transceiver. To achieve this end, latches are eliminated from the 2-to-1 MUX and the 0.5UI phase shift between the input 28Gb/s signals is implemented by using the four-phase 14GHz clock that drives the preceding 32-to-2 MUXs. In this design, the timing between input data and the clock signal for the 2-to-1 MUX is adjusted by adjusting the phase of the 14GHz clock signals output from the frequency divider. Also, the 2-to-1 MUXs are kept small to reduce the 28GHz clock load and use a four-stage pre-driver to drive the SST output stages. The pre-driver consists of four inverter stages with series peaking inductors placed in between the stages to enhance the bandwidth. This configuration achieves better bandwidth and power efficiency compared to using a bigger 2-to-1 MUX and fewer stages of pre-driver.

Due to the limitation of the measurement equipment, only a single lane was measured, although the transceiver has a two-lane configuration. The 56.2Gb/s transmitter output eye observed through a channel loss of 7.2dB at 28.1GHz is shown in Fig. 3.5.5(a). The gain boost of the 2-tap FFE is 4dB at half the baud frequency. The total jitter is 8.8ps for a BER of 10^{-15}, with an rms random jitter of 288fs and duty-cycle-distortion of 125fs for a 56.2Gb/s PRBS15 data pattern at room temperature. The jitter tolerance curve measured using a 56.2Gb/s PRBS9 and PRBS15 pattern at BER=10^{-12} satisfies an OIF CEI-56G-VSR MASK (Fig. 3.5.5(b)). The RX internal eye diagram for PRBS15 pattern at BER=10^{-9} are measured using an on-chip eye monitor (Fig. 3.5.5(c)). In the jitter tolerance measurement, the channel loss excluding the socket and package loss, which is estimated to be about 5dB, was 18.4dB at 28.1GHz. The performance summary is shown in Fig. 3.5.6 and a photograph of the test chip fabricated in 28nm CMOS is shown in Fig. 3.5.7. The transceiver achieved 18.4dB loss compensation at 56.2Gb/s. Power consumption of the transceiver was 247mW/lane from a 0.96V supply. The area of the 2-lane transceiver was 1.0×1.4mm^2.

References:

[1] T. Shibasaki et al., "A 56-Gb/s Receiver Front-End with a CTLE and 1-Tap DFE in 20-nm CMOS," *IEEE Symp. VLSI Circuits*, pp. 112-113, Jun. 2014.
[2] J. Lee et al., "56Gb/s PAM4 and NRZ SerDes Transceivers in 40nm CMOS," *IEEE Symp. VLSI Circuits*, pp. 118-119, Jun. 2015.
[3] J. Han et al., "A 60Gb/s 173mW Receiver Frontend in 65nm CMOS Technology," *IEEE Symp. VLSI Circuits*, pp. 230-231, Jun. 2015.
[4] S. Parikh et al., "A 32Gb/s Wireline Receiver with a Low-Frequency Equalizer, CTLE and 2-Tap DFE in 28nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 28–29, Feb. 2013.
[5] H. Miyaoka et al., "A 28-Gb/s 4.5-pJ/bit Transceiver With 1-Tap Decision Feedback Equalizer in 28-nm CMOS," *Asian Solid-State Circuits Conf.,* pp. 245-248, Nov. 2015.

 978-1-4673-9467-3/16/$31.00

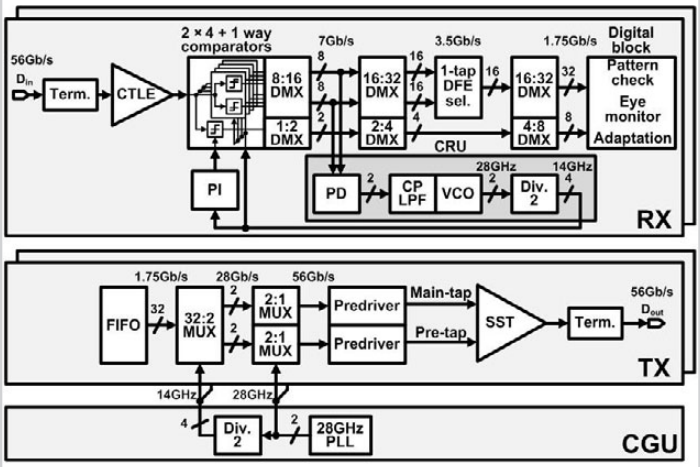

Figure 3.5.1: 56Gb/s transceiver block diagram.

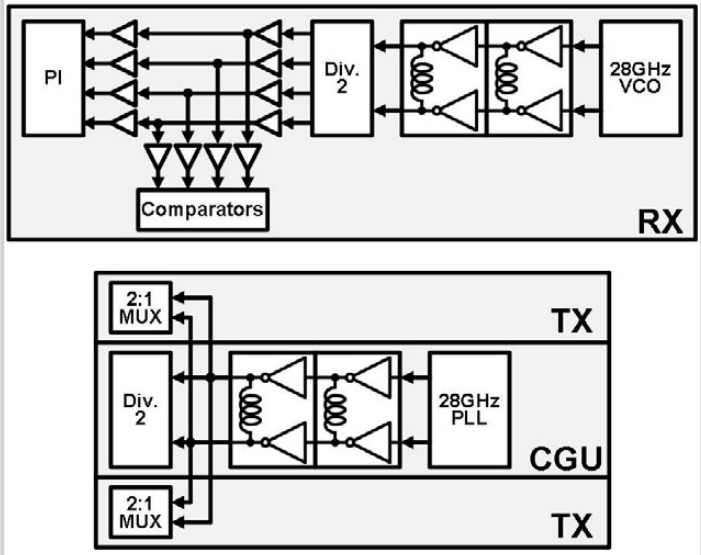

Figure 3.5.2: Transceiver clock-path design.

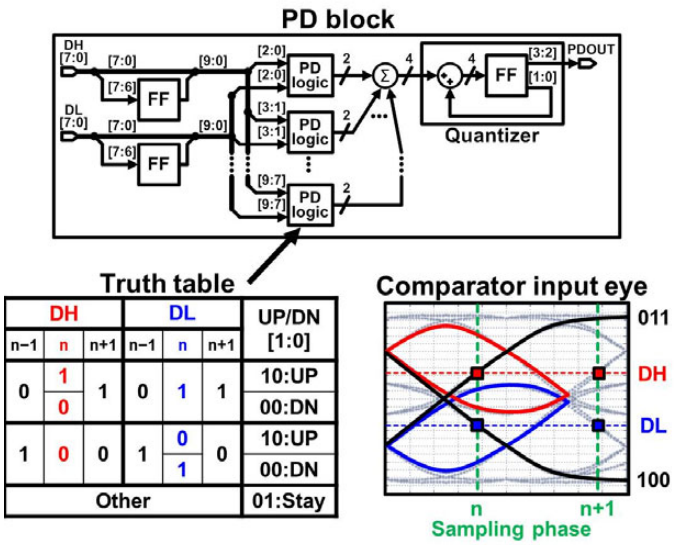

Figure 3.5.3: PD block diagram and truth table.

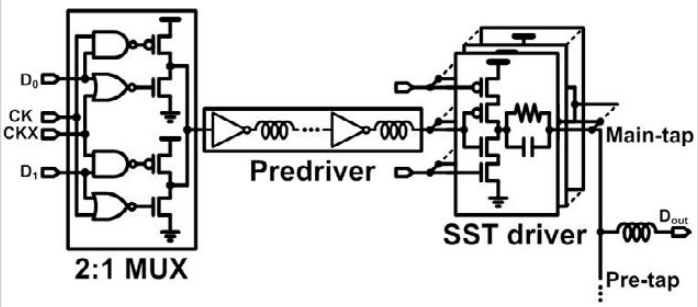

Figure 3.5.4: Transmitter 2-to-1 MUX and SST driver stage.

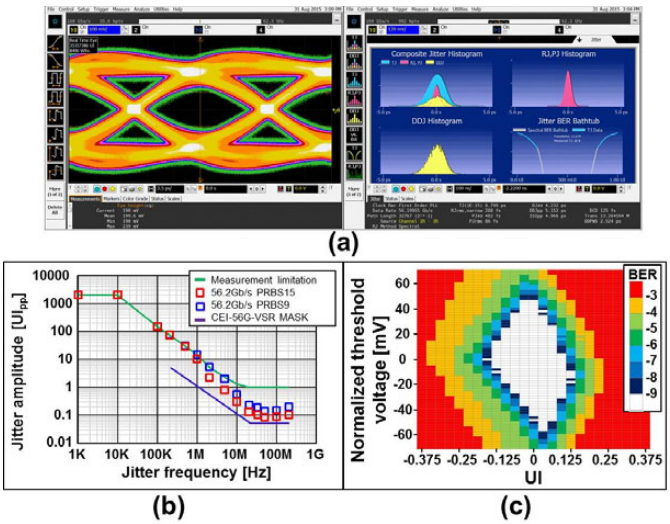

Figure 3.5.5: (a) TX eye diagram, (b) RX jitter-tolerance curve and (c) RX internal eye diagram.

	[1]	[2]	[3]	This work
Technology	20nm CMOS	40nm CMOS	65nm CMOS	28nm CMOS
Supply voltage [V]	0.9	1.2	1.2/1.0	0.96
Data rate [Gb/s]	56	55.5-56.5	60	56.2
Channel loss [dB]	23	–	–	18.4
Power [mW/Gb/s]	RX: 3.16	TX: 8.00 RX: 3.93	RX: 2.88	TX: 1.87 RX: 2.53
Area [mm^2]	0.34	TX: 2.10 RX: 1.10	RX: 0.16	1.40
TX function	–	4:1 MUX 60GHz PLL PRBS	–	32:1 MUX 2-tap FFE 28GHz PLL
RX function	CTLE 1-tap DFE 4:16 DMX	CTLE CDR 1:8 DMX	CTLE 2-tap FFE 3-tap DFE	CTLE 1-tap DFE CDR 4:32 DMX Eye mon. Adaptation

Figure 3.5.6: Performance summary.

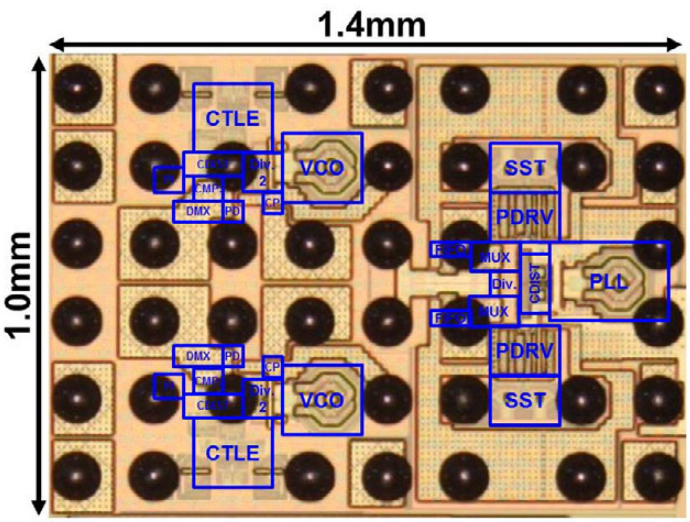

Figure 3.5.7: Chip photograph.

978-1-4673-9467-3/16/$31.00

3.6 A 45Gb/s PAM-4 Transmitter Delivering 1.3V_{ppd} Output Swing with 1V Supply in 28nm CMOS FDSOI

Matteo Bassi[1], Francesco Radice[2], Melchiorre Bruccoleri[2], Simone Erba[3], Andrea Mazzanti[1]

[1]University of Pavia, Pavia, Italy, [2]STMicroelectronics, Cornaredo, Italy, [3]STMicroelectronics, Pavia, Italy

The development of next-generation electrical link technology to support 400Gb/s standards is underway [1-5]. Physical constraints paired to the small area available to dissipate heat, impose limits to the maximum number of serial interfaces and therefore their minimum speed. As such, aggregation of currently available 25Gb/s systems is not an option, and the migration path requires serial interfaces to operate at increased rates. According to CEI-56G and IEEE P802.3bs emerging standards, PAM-4 signaling paired to forward error correction (FEC) schemes is enabling several interconnect applications and low-loss profiles [1]. Since the amplitude of each eye is reduced by a factor of 3, while noise power is only halved, a high transmitter (TX) output amplitude is key to preserve high SNR. However, compared to NRZ, the design of a PAM-4 TX is challenged by tight linearity constraints, required to minimize the amplitude distortion among the 4 levels [1]. In principle, current-mode (CM) drivers can deliver a differential peak-to-peak swing up to 4/3(V_{DD}-V_{OV}), but they struggle to generate high-swing PAM-4 levels with the required linearity. This is confirmed by recently published CM PAM-4 drivers, showing limited output swings even with V_{DD} raised to 1.5V [2-4]. Source-series terminated (SST) drivers naturally feature better linearity and represent a valid alternative, but the maximum differential peak-to-peak swing is bounded to V_{DD} only. In [5], a dual-mode SST driver supporting NRZ/PAM-4 was presented, but without FFE for PAM-4 mode. In this paper, we present a PAM-4 transmitter leveraging a hybrid combination of SST and CM driver. The CM part enhances the output swing by 30% beyond the theoretical limit of a conventional SST implementation, while being calibrated to maintain the desired linearity level. A 5b 4-tap FIR filter, where equalization tuning can be controlled independently from output matching, is also embedded. The transmitter, implemented in 28nm CMOS FDSOI, incorporates a half-rate serializer, duty-cycle correction (DCC), >>2kV HBM ESD diodes, and delivers a full swing of 1.3V_{ppd} at 45Gb/s while drawing 120mA from a 1V supply. The power efficiency is ~2 times better than those compared in this paper.

Figure 3.6.1(a) illustrates the core of the proposed output driver. Without current sources (I_s/3, 2/3I_s), it consists of CMOS inverters switching resistors 3R_L and 3/2R_L to either V_{DD} or ground. Branch impedances are properly scaled to generate the four PAM-4 levels and to ensure output matching to R_L=50Ω. Two issues affect the classic SST driver topology. First, because the device on-resistance of the inverters is nonlinear, it needs to be small compared to the explicit resistors, requiring large inverter sizes. This leads to higher pre-driver loading and sets a trade-off between TX linearity and power consumption. Second, the impedance match condition constrains the maximum differential peak-to-peak swing $V_{o\text{-}ppd}$ to V_{DD}. Rising the supply to increase $V_{o\text{-}ppd}$ needs power-hungry level translators and thick oxide devices or protection circuits, penalizing TX speed or efficiency.

The current sources (I_s/3, 2/3I_s) switched by differential pairs in Fig. 3.6.1(a) are introduced to tackle the aforementioned issues. By injecting additional currents to the output nodes, the output swing is raised to V_{DD}+2$I_s R_L$ [6]. With V_{DD}=1V and I_s=3mA, $V_{o\text{-}ppd}$ is boosted from 1V to 1.3V. Note that compared to increasing V_{DD} to 1.3V, this solution gains 30% efficiency, and does not require level translators. To compensate for the nonlinear on-resistance of the small inverters, and ensure high linearity even with small voltage headroom across the current sources, the latter are calibrated in a master-slave fashion with the scaled replicas of the driver shown Fig. 3.6.1(b). Four DACs generate the nominal PAM-4 output levels (V_{00}, V_{01}, V_{10} and V_{11}), and op-amps closed in coupled feedback loops tune the current sources.

Figure 3.6.2 shows the architecture of the FIR equalizer leveraging the modified SST driver. It consists of 4 symbol-spaced taps, where, to save area and power, a bank of 2:1 multiplexers switch either the cursor or the pre/post-cursor streams to the driver. This choice sets the maximum weight of the pre- and post-cursors taps to ¼, while multiplexing inverted-cursor and cursor in the main tap allows for a fine tuning of the output amplitude. Note that compared to CM drivers, where coefficient pre-allocation limits cursor size when FIR is enabled, the main tap here is always automatically maximized. This feature is attractive to deliver the highest PAM-4 swing to the output. Coefficients c_i[0:5] are routed to the 6 cells connected in parallel which build a tap. To ensure monotonicity of the tap gains against device mismatches, the cells are scaled by a factor 1/k=[1,2,4,8,8,8], corresponding to control words encoded with 4 binary and 2 thermometric bits.

The 2:1 half-rate serializer shown in Fig. 3.6.3 provides full-rate symbol-spaced data to drive the FIR taps of the output stage. Only for testing purposes, it is fed by two external full-rate NRZ streams (MSBin and LSBin), split into half-rate sequences on-chip. Input buffers convert MSBin and LSBin data to CMOS levels and route it to 4 serializer slices (MSBp, MSBn, LSBp and LSBn) driven by half-rate clock CLKin, which is fed through a buffer including a DCC circuit. The full-rate data is first split into two half-rate sequences. Even/odd data are delayed by latches and then re-multiplexed to 16 full-rate streams. Delays, ΔT, are inserted to match sample timing. A matrix of switches is employed as a sign selector before delivering data to the driver core. Up to 5 FIR slices, identical to the one shown in Fig. 3.6.2, can be connected in parallel to match the output impedance to R_L. The advantage of this architecture is that the equalization setting can be maintained independently from the matching condition [7]. ESD protections are critical elements for high data rate operation. Two 200V MM / 500V CDM / >>2kV HBM compliant ESDs are included at the output of the TX, accounting for a parasitic capacitance of ~250fF each. To improve output return loss and enable high data rates, asymmetric T-coils are employed.

The PAM-4 TX is implemented in a 28nm CMOS FDSOI process by STMicroelectronics. Test-chips (micrograph in Fig. 3.6.7) have been encapsulated on plastic flip-chip BGA packages and mounted on PCBs for measurements. An Anritsu MP1800A unit with two independent PPGs provides MSB/LSB data streams and clock, and the output is connected to a sampling scope. The channel profile of the measurement setup, including cable, PCB trace and package losses is show in Fig. 3.6.4. The effectiveness of the proposed solution to increase the TX output swing is first demonstrated at 10Gb/s, where TX FIR filtering is not needed. Enabling the current injection enhances the output swing from 1V_{ppd} to 1.3V_{ppd}, thus providing a 30% increase in the full scale output amplitude.

The TX FFE provides a maximum boost of 6dB, compensating channel losses up to ~12GHz, as shown in Fig. 3.6.4. In this condition, TX is tested at the data rate of 45Gb/s, and results are shown in Fig. 3.6.5. With current injection turned off, the TX delivers an output swing of ~550mV_{ppd}. When turned on, output swing is enhanced by 28%, raising to ~700mV_{ppd}. At this rate, the driver draws 120mA from a 1V supply.

Increasing output swing while minimizing eye distortion is the target of this work. As proposed by emerging standards, distortion is measured using a special sequence of 10 voltage levels, each of 16UI duration [1]. Figure 3.6.5 shows a measurement of the resulting set of PAM-4 symbols at 45Gb/s. At this rate, the level separation mismatch ratio R_{LM}=3min(V_D-V_C,V_C-V_B,V_B-V_A)/(V_D-V_A) is calculated for several chip samples and is better than the spec (under discussion) of 0.92 with margin and independent from FFE coefficients settings.

Measured results are summarized and compared with published PAM-4 transmitters in the table of Fig. 3.6.6. Thanks to the use of a hybrid mostly-SST architecture, the proposed TX is able to deliver the highest reported output swing with ~2× better power efficiency than those compared in Fig. 3.6.6. These features, paired with good linearity, are particularly effective to accommodate the demand for high output swing, linearity and efficiency required by emerging PAM-4 interfaces.

References:

[1] Optical Interconnect Forum [Online]. http://www.oiforum.com

[2] A. Nazemi et al., "A 36Gb/s PAM4 transmitter using an 8b 18GS/S DAC in 28nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 58-59 Feb. 2015.

[3] P. C. Chiang et al., "60Gb/s NRZ and PAM4 transmitters for 400GbE in 65nm CMOS," *ISSCC Dig. Tech. Papers*, pp.42-43, Feb. 2014.

[4] C. Menolfi et al., "A 25Gb/s PAM4 transmitter in 90nm CMOS SOI," *ISSCC Dig. Tech. Papers*, pp.72-73, Feb. 2005.

[5] J. Kim et al., "A 16-to-40Gb/s quarter-rate NRZ/PAM4 dual-mode transmitter in 14nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 60-61, Feb. 2015

[6] A. Joy et al., "Analog-DFE-Based 16Gb/s SerDes in 40nm CMOS That Operates Across 34dB Loss Channels at Nyquist with a Baud Rate CDR and 1.2Vpp Voltage-Mode Driver," *ISSCC Dig. Tech Papers*, pp.350-351, Feb. 2011.

[7] M. Kossel et al., "A T-Coil-Enhanced 8.5Gb/s High-Swing SST Transmitter in 65nm Bulk CMOS With <-16dB Return Loss Over 10GHz Bandwidth," *IEEE J. Solid-State Circuits*, vol.43, no.12, pp.2905-2920, Dec. 2008.

 978-1-4673-9467-3/16/$31.00

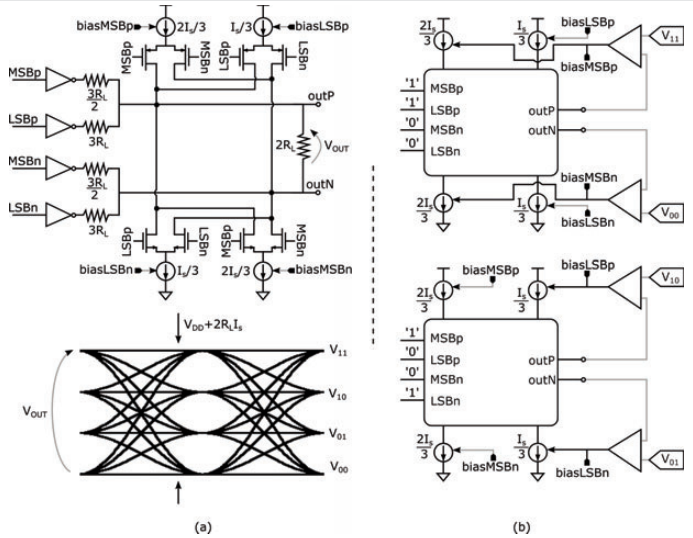

Figure 3.6.1: (a) PAM-4 SST driver with current sources ($I_s/3$, $2/3I_s$) to enhance output swing and (b) TX replicas to calibrate the PAM-4 output voltage levels.

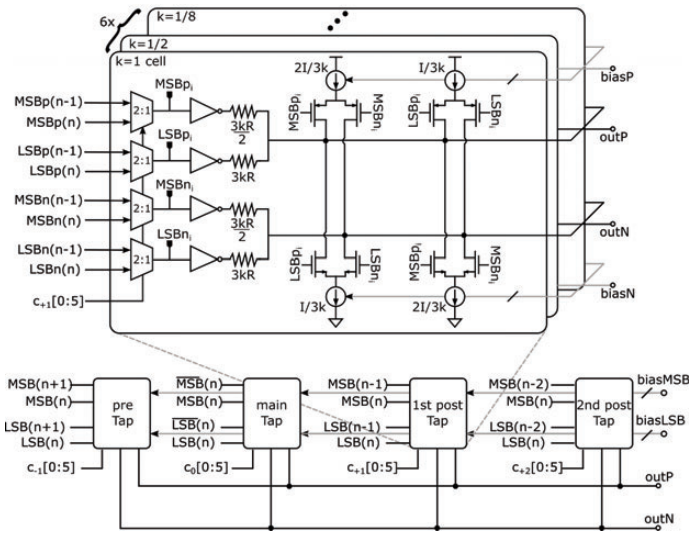

Figure 3.6.2: Block diagram of one slice of the SST driver.

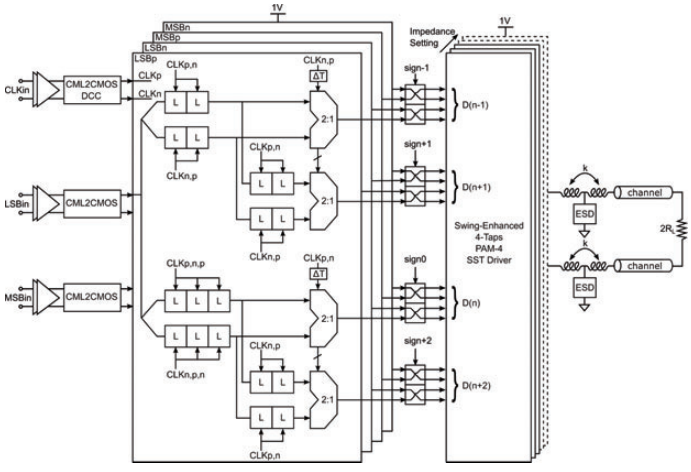

Figure 3.6.3: Block diagram of the serializer, SST driver and output network.

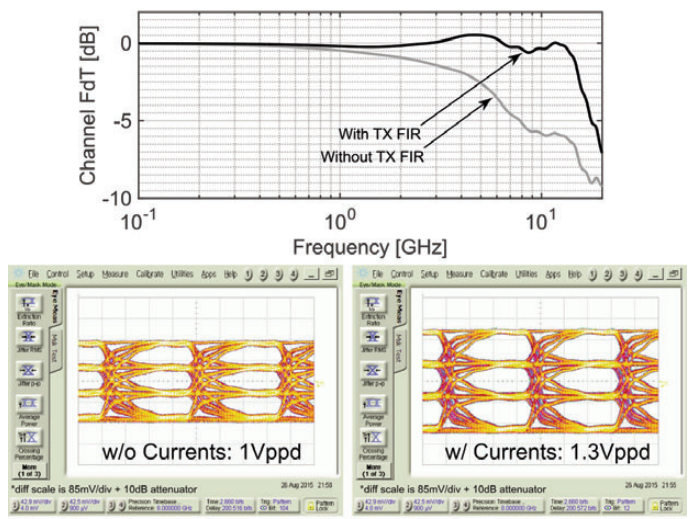

Figure 3.6.4: (top) measured test channel profile and (bottom) measured TX output eyes at 10Gb/s w/ and w/o current sources.

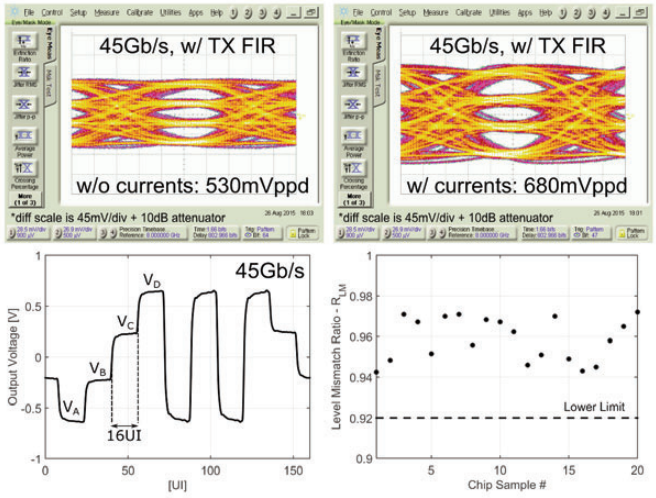

Figure 3.6.5: (top) measured TX output eyes at 45Gb/s including test channel w/ and w/o current sources and (bottom) measured PAM-4 eye distortion at 45Gb/s for several chip samples.

Ref.	[4]	[2]	[3]	[5]	This Work
CMOS Technology	90nm SOI	28nm	65nm	14nm	28 nm FDSOI
Driver Topology	CML	CML	CML	SST	SST
TX-FFE	4-taps	DAC	3-taps	No	4-taps
ESD	Yes	No	No	Yes	Yes
Data-Rate [Gb/s]	25	36	60	40	45
Output Swing V_{out} [Vppd]	0.84[1]	0.8	0.250	0.9[3]	1.3
Vdd [V]	1	1.5	1.2	N/A	1
Power P_{DC} [mW]	102	84[2]	205[2]	167.5[2]	120
Power Efficiency ($V_{out}^2/2R)/P_{DC}$ [%]	3.4	3.8	0.15	2.42	7
mw/Gbps	4	2.33	3.4	4.18	2.6
Area [mm²]	0.052	0.05	1.14	0.0279	0.28

1 Amplitude from picture. Loss recovered by FFE de-embedded.
2 Not including PLL and clock distribution power.
3 Amplitude from picture. Loss recovered by software CTLE de-embedded.

Figure 3.6.6: TX performance comparison.

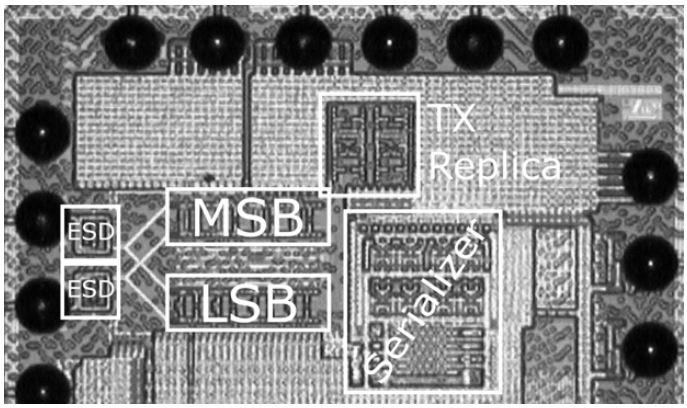

Figure 3.6.7: Die micrograph.

978-1-4673-9467-3/16/$31.00

3.7 A 40-to-64Gb/s NRZ Transmitter with Supply-Regulated Front-End in 16nm FinFET

Yohan Frans, Scott McLeod, Hiva Hedayati, Mohamed Elzeftawi, Jin Namkoong, Winson Lin, Jay Im, Parag Upadhyaya, Ken Chang

Xilinx, San Jose, CA

Due to increasing bandwidth demand in data centers and telecommunication infrastructures, the maximum data-rate of wireline transceivers is projected to double from 32Gb/s to 64Gb/s while keeping the same power envelope. This paper presents the design of a 64Gb/s NRZ transmitter for short-reach electrical links in a 16nm FinFET process. It is applicable to standards such as CEI-56-VSR/MR, using power efficient techniques that take into considerations FinFET device properties: low DIBL, negligible body effect, low junction capacitance, low channel leakage, high intrinsic gain, high gate capacitance and resistance, high flicker noise, and steep CV curve in accumulation region.

Quad-rate architecture (Fig. 3.7.1) is chosen to relax serialization timing [1]. The transmitter uses 3-tap pre-emphasis (one pre, main and post tap) to equalize the channel loss. The transmitter front-end consists of four identical slices (4:1 MUX and driver in each slice): two main-tap slices and one slice each for programmable pre/main and post/main taps. The driver slices are open-drain CML stages whose outputs are shorted together and connected to a T-coil-enhanced termination network. Four-phase quad-rate CML clock inputs are converted to CMOS, buffered, and distributed to the front-end slices. The phases of the four-phase clocks are calibrated.

In the previous 4:1 MUX implementations [1], a part of data qualification is performed using low-fanout (required to meet bandwidth) dynamic logic and another part is performed at the final multiplexing stage. This approach is susceptible to inter-symbol interference, especially when the slew rate of the multiplexer inputs is not fast enough [2]. Several FinFET device properties (lower junction capacitance, lower DIBL, and negligible body effect) significantly improve the performance of low-fanout stacked logic gates (Fig. 3.7.2) beyond geometric scaling. The 4:1 MUX implemented in this design takes advantage of this by performing the entire data qualification before the multiplexing stage using static CMOS gates [3]. It consists of four CMOS pulse generators and a tailless CML multiplexing stage. Each pulse generator sequentially ANDs one quad-rate data (d0–d3) with two adjacent four-phase clocks (clk0–clk3). The pseudo-differential CMOS outputs of the pulse generator (z0–z3 and their complements) are both driven to 0V for 3UIs and complementary for 1UI. The four pulse generators thus generate four 1UI pulses, each pulse staggered by 1UI. The multiplexer converts the 1UI CMOS pulses into a differential full-rate bit-stream (V_o) at proper common-mode level amenable to the subsequent main CML driver. A tailless CML structure reduces the size of NMOS devices in the multiplexer, resulting in lower capacitance at both the input and output of the multiplexer. Small peaking inductors are placed at the multiplexer output, which is predominantly self-loaded, to provide moderate bandwidth extension of ~1.5×. A PMOS load is used since it has lower total parasitic capacitance (including device and wiring) compared to a passive resistor. The subsequent CML driver (with tail current source) suppresses common mode ripples present at the multiplexer output due to its tailless structure, thus reducing transmitter output noise caused by common-mode-to-differential conversion at the package and PCB.

The clock buffers and pulse-generators are built in CMOS to ensure a fast edge-rate, which minimizes its random jitter. An on-chip regulator (Fig. 3.7.3) provides a clean supply to these circuits. The pseudo-differential nature of the pulse generators ensures that the transient current is independent of the data pattern. The resulting transient current exhibits periodic droops at the full-rate frequency that can be effectively suppressed by a de-coupling cap. Moreover, the residual voltage ripple has no impact on clock jitter because it is exactly at full-rate frequency. By nominally setting the regulated supply (VREG) level equal to the unregulated supply (AVCC), setup/hold timing constraints between the 16:4 serializer and the 4:1 MUX can be met by simply adjusting clock-tree delays during the design phase. The regulator [4] consists of a high-gain, low-bandwidth amplifier, followed by a thin-oxide NMOS source follower that provides >20dB PSRR thanks to high intrinsic gain in the FinFET process. A 1.8V auxiliary supply VCCAUX (shared with the PLL) is used to supply a small amount of power to the amplifier in order to provide enough overdrive voltage to the thin-oxide NMOS.

Clock phase errors are sensed at the leaf nodes of the regulated CMOS clock distribution to minimize residual error (Fig 3.7.4). Since clk0 (clk1) and clk2 (clk3) are differentially coupled using cross-coupled inverters along the CMOS clock chain, any deviation from the ideal 180-degree phase shift will manifest as duty-cycle distortion (DCD). The DCD is detected by checking differences in common-mode level between clk0 (clk1) and clk2 (clk3). DCD corrections up to ±3ps with 100fs resolution are performed in the CMOS domain by digitally adjusting pull-up and pull-down strengths of inverters. This is accomplished by having a bank of small tri-state buffers in parallel with the main inverter buffer, where each pull-up and pull-down leg is controlled independently. Quadrature phase error between the clk0–clk2 pair and the clk1–clk3 pair is detected using symmetric CML XOR gates [5]. High intrinsic gain and low channel leakage in the 16nm FinFET makes it possible to build a high-performance offset-canceling (auto-zero) comparator, resulting in <250fs detection resolution. Quadrature correction up to ±3ps with 100fs resolution is performed in the CML domain by interpolating the outputs of two CML buffers separated by an RC delay [5].

The LC PLL, shown in Fig. 3.7.5, uses two LC-VCOs – only one selected at a time – to provide the required tuning range. The LC-VCOs oscillate at half-rate (e.g. 32GHz for 64Gb/s operation) and are followed by a tailless CML quadrature divider to generate four-phase quad-rate clocks. In this 16nm FinFET process, the varactor CV curve is much steeper compared to the planar process used in [5] and it varies significantly across PVT. The flicker noise is also much higher in the 16nm FinFET compared to the planar process. The LC-VCO circuit is designed to mitigate these issues. A PMOS-only topology sets the VCO output common-mode level to ground through the low series resistance of the inductors. The low-impedance ground-level common-mode alleviates modulation due to the flicker noise of the bias current source. The LC-VCO is designed to have a large output swing (>800mV_{ppd}) to help linearize the steep varactor CV curve. AC-coupling capacitors are placed between the varactor and LC-tank terminals and the gate terminal of the varactors is connected to an adjustable bias voltage (v_{bias}) to set the varactor to operate in the optimum region. Common wireline transceiver systems operating above 40Gb/s are expected to have a CDR bandwidth above 15MHz. Therefore, the LC-VCO design is optimized to achieve low phase noise above 10MHz.

Figure 3.7.6 shows the transmitter output captured using Keysight DSA-X96204Q oscilloscope after a 20mm package trace, 1.5" PCB trace, connector, 18" cable, and 50GHz DC-blocking capacitors. The scope measures 150fs_{rms} RJ, 223fs_{pp} PJ, and 30fs_{pp} DCD with a 64Gb/s clock pattern. It measures 600fs_{pp} PJ and 2.8ps ISI with a 64Gb/s PRBS-7 pattern. The clock-pattern measurements over 39 units show significantly tighter PJ distribution when phase error correction is turned on, indicating the effectiveness of the phase error correction. In the worst-case part, the phase error correction reduces clock-pattern PJ_{pp} from 1.8ps to 400fs. The overall jitter performance is significantly better than other published >60Gb/s transmitters [6, 7].

Figure 3.7.7 shows the die micrograph and performance summary of the transmitter. Two transmitter channels and one shared PLL are fabricated. The transmitter consumes 225mW per channel. The amortized PLL, clock distribution, and bias generation power per channel is 115mW.

Acknowledgment:
The authors would like to thank Xilinx Serdes team for their contributions in circuit design, chip bring-up, and measurements.

References:
[1] A. Hafez et al., "A 32-to-48Gb/s Serializing Transmitter Using Multiphase Sampling in 65nm CMOS", *ISSCC Dig. Tech. Papers*, pp. 38-39, Feb. 2013.
[2] C.-K. K. Yang, "Design of High-Speed Serial Links in CMOS," Ph. D. thesis, University of California Los Angeles, CA, USA, 1998.
[3] A. Fiedler et al, "A 1.0625 Gbps transceiver with 2x-oversampling and transmit signal pre-emphasis," *ISSCC Dig. Tech. Papers*, pp. 238-239, Feb. 1997.
[4] E. Alon, "Measurement and regulation of on-chip power supply noise," Ph. D. thesis, Stanford University, Stanford, CA, USA, 2006.
[5] P. Upadhyaya et al., "3.3 A 0.5-to-32.75Gb/s flexible-reach wireline transceiver in 20nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 56-57, 22-26 Feb. 2015.
[6] M.-S. Chen et al., "A 50 - 64 Gb/s Serializing Transmitter With a 4-tap, LC-Ladder-Filter-Based FFE in 65-nm CMOS", *IEEE Custom Integrated Circuits Conf.*, pp. 1-4, Sept. 2014.
[7] P. Jiang et al., "60Gb/s NRZ and PAM4 Transmitters for 400GbE in 65nm CMOS", *ISSC Dig. Tech. Papers*, pp. 42-43, Feb. 2014.

 978-1-4673-9467-3/16/$31.00

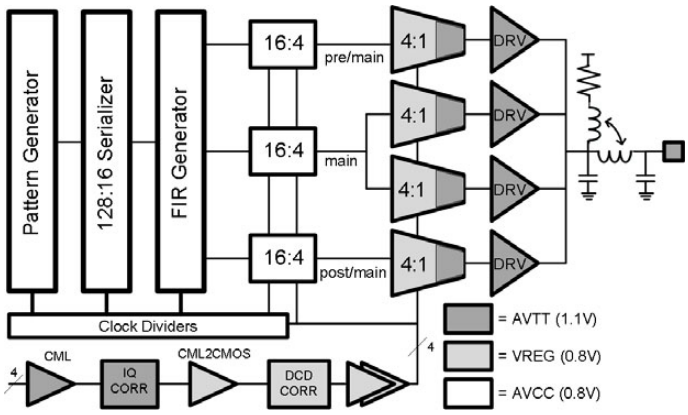

Figure 3.7.1: Transmitter block diagram.

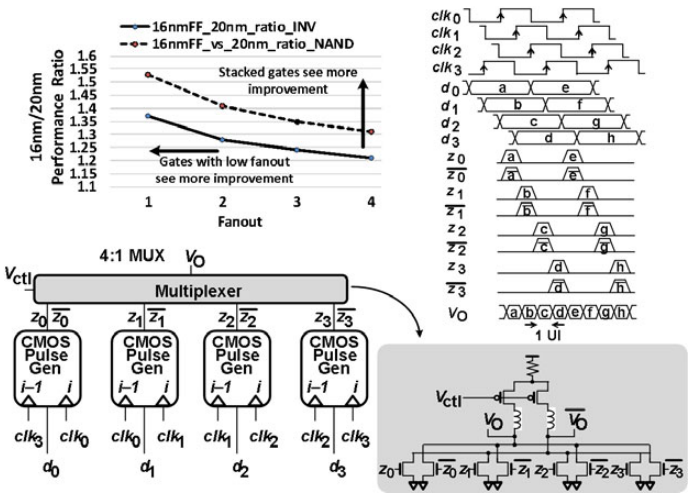

Figure 3.7.2: 4:1 MUX implementation and timing diagram.

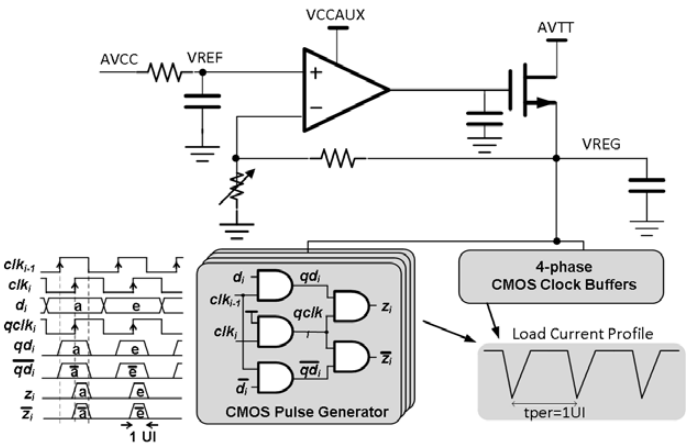

Figure 3.7.3: Voltage regulator with priodic CMOS load current profile.

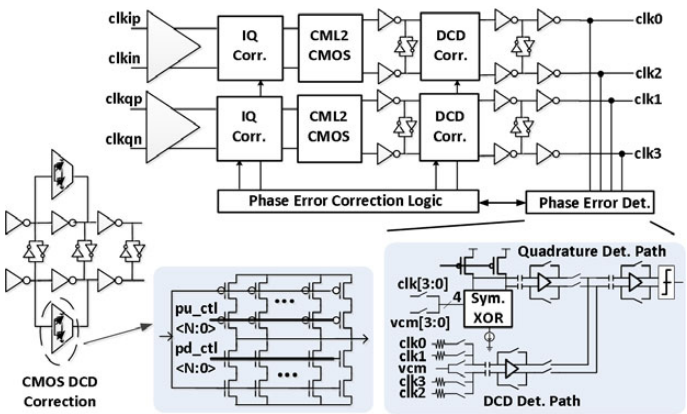

Figure 3.7.4: Phase error correction.

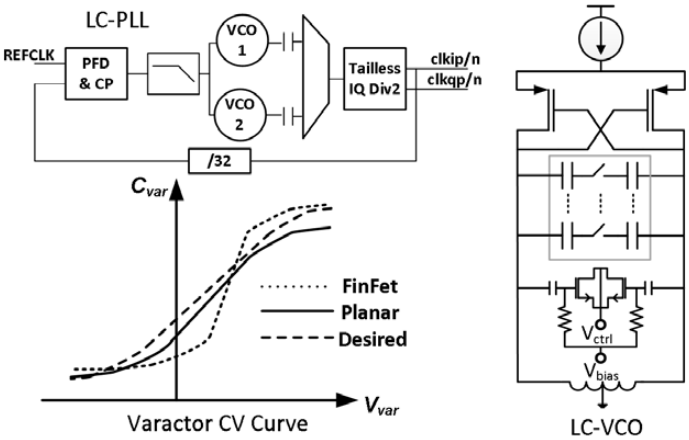

Figure 3.7.5: LC-VCO circuit.

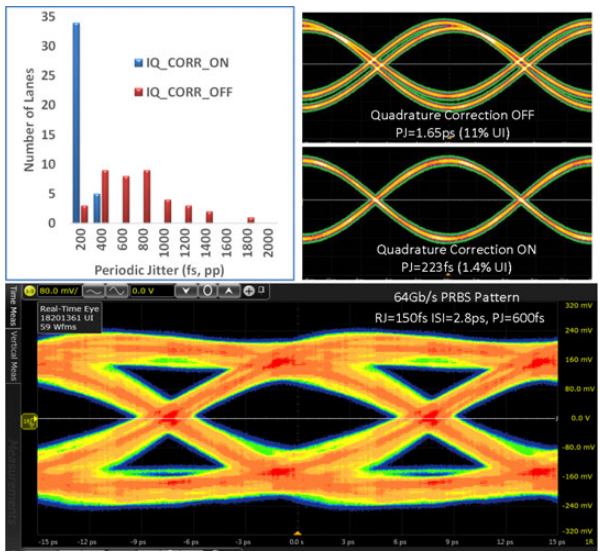

Figure 3.7.6: 64Gb/s clock pattern and PRBS-7 pattern outputs.

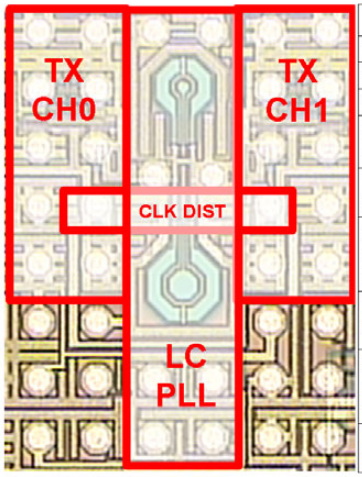

Performance Summary	
Technology	16nm FinFET
Data-rate	40-64 Gb/s
FFE	3-tap
Area (TX Only)	0.32mm2 (C4 limited)
TX Power per channel	225mW @ 64Gb/s
PLL & Clock Distribution Power per channel	115mW @ 64Gb/s
TX Amplitude (No EQ, diff-pp)	800mV
Wide-band TX Jitter (Clock Pattern) @ 64Gb/s	RJ-rms=150fs PJ-pp=200fs DCD=30fs
TX Jitter (PRBS) @ 64Gb/s	PJ=600fs ISI=2.8ps

Figure 3.7.7: Performance summary and die micrograph.

978-1-4673-9467-3/16/$31.00